

An Information Roadmap Method for Robotic Sensor Path Planning

G. Zhang · S. Ferrari · M. Qian

Received: 16 May 2008 / Accepted: 11 February 2009 / Published online: 11 March 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract A new probabilistic roadmap method is presented for planning the path of a robotic sensor deployed in order to classify multiple fixed targets located in an obstacle-populated workspace. Existing roadmap methods have been successful at planning a robot path for the purpose of moving from an initial to a final configuration in a workspace by a minimum distance. But they are not directly applicable to robots whose primary objective is to gather target information with an on-board sensor. In this paper, a novel information roadmap method is developed in which obstacles, targets, sensor's platform and field-of-view are represented as closed and bounded subsets of an Euclidean workspace. The information roadmap is sampled from a normalized information theoretic function that favors samples with a high expected value of information in configuration space. The method is applied to a landmine classification problem to plan the path of a robotic ground-penetrating radar, based on prior remote measurements and other geospatial data. Experiments show that paths obtained from the information roadmap exhibit a classification efficiency several times higher than that of existing search strategies. Also, the information roadmap can be used to deploy non-overpass capable robots that must avoid targets as well as obstacles.

Keywords Information-driven sensor planning · Path planning · Robot sensing · Sensor fusion · Probabilistic roadmap methods · Demining · Information entropy

G. Zhang · S. Ferrari (✉)
Department of Mechanical Engineering and Materials Science,
Duke University, Durham, NC 27708, USA
e-mail: sferrari@duke.edu

S. Ferrari · M. Qian
Department of Electrical and Computer Engineering,
Duke University, Durham, NC 27708, USA

1 Introduction

Sensor planning refers to the problem of determining a strategy for gathering sensor measurements to support a sensing objective, such as target classification. When the sensors are installed on robotic platforms an important part of the problem is planning the sensor path [1–3]. Several approaches have been proposed for planning the path of mobile robots with on-board sensors to enable navigation and obstacle avoidance in unstructured dynamic environments, e.g., [4–10]. However, these methods are not directly applicable to robotic sensors whose primary objective is to support a sensing objective, rather than to simply navigate a dynamic environment [11]. The reason is that these methods focus on how the sensor measurements can best support the robot motion, rather than focusing on the robot motions that best support the sensing objective [11]. This paper focuses on the latter, addressing the problem of planning the path of a robotic sensor in order to classify multiple fixed targets distributed in an obstacle-populated workspace, which arises in many applications, such as, robotic mine hunting [12], cleaning [1], and monitoring of urban environments [13], manufacturing plants [14], and endangered species [15].

Existing approaches to sensor path planning include coverage path-planning [2, 11], random [2], grid [16], and optimal search strategies [16, 17]. Optimal search strategies typically outperform other approaches in applications where a-priori information is available, such as sensor models, environmental conditions, and prior measurements [16]. However, they do not yet provide a systematic and general approach for sensor path planning in geometric sensing problems. *Geometric sensing* problems require a description of the geometry and position of the targets and of the sensor's field-of-view [18]. Viewpoint planning has been shown by several authors to be an effective approach for optimally placing or moving vision sensors based on the target geometry and sensor field-of-view, using weighted functions or tessellated space approaches [18–20]. Probabilistic deployment computes a search path based on the probability of finding a target in every unit bin of a discretized, obstacle-free workspace [1, 21, 22]. This paper presents a new information roadmap method that combines information-driven sensor planning with probabilistic roadmap methods (PRMs) to compute the sensor *path* in a geometric sensing problem. Its advantages over existing techniques are that it takes into account the motion and geometry of closed and bounded subsets of an Euclidian space, representing the sensor's platform and field-of-view, as well as the geometry and position of multiple fixed targets and obstacles in the workspace.

Traditionally, PRMs have been used to plan the motions of a robot with geometry \mathcal{A} , in order to avoid collisions with multiple fixed obstacles in a workspace \mathcal{W} [23–26]. While robot path planning typically aims to optimize a deterministic additive function such as Euclidian distance, sensor planning aims to optimize a stochastic sensing objective that is not necessarily additive. Another basic difficulty in sensor planning is that, although the measurements ultimately determine the sensor performance, they cannot be factored into the planning problem because the sensor's position must be planned prior to obtaining the sensor's measurements [27–31]. Recently, several authors have shown that this difficulty can be overcome by an approach known as information-driven sensor planning, which uses information theoretic functions to estimate the measurements' value prior to deploying the sensor [27–30].

Existing PRMs and sensor modeling approaches used for information-driven sensor planning are reviewed in Section 3. Existing methods, however, are not directly applicable to the geometric sensor path planning problem formulated in Section 2, because they do not account for the geometry of the sensor and the workspace. In Section 4, a novel information roadmap method is developed that combines PRM and information-driven sensor planning by sampling a normalized information theoretic function defined over the robot's free configuration space. The information theoretic function is the expected conditional entropy reduction (EER) of potential targets, presented in Section 4.1. A new hybrid sampling strategy is presented in Section 4.2 to generate an information roadmap that contains a high density of milestones with high EER, as well as milestones that capture the connectivity of the free configuration space. Section 4.3 presents a new query phase for searching the information roadmap using an A^* -type algorithm that returns the path of maximum information profit whenever one exists, and returns failure otherwise.

The information roadmap method is demonstrated by planning the path of a mine hunting ground-penetrating radar (GPR) installed on a ground robot (Section 5). The results in Section 6 show that the method outperforms existing sensor path-planning methods that are applicable to geometric sensing (namely, complete coverage [2, 11] and random search [2]) under a wide range of workspace conditions and geometries, increasing the average classification efficiency by up to one order of magnitude. Also, the information roadmap can be used to deploy non-overpass capable robotic sensors that may be damaged or destroyed when driving over landmines. As a result, the robotic sensor avoids collisions with potential mines, while obtaining measurements from them, both in wide-open regions of the workspace and inside narrow passages.

2 Problem Formulation and Assumptions

This paper addresses the problem of planning the path of a robotic sensor with a platform geometry $\mathcal{A} \subset \mathbb{R}^2$, and a field-of-view geometry $\mathcal{S} \subset \mathbb{R}^2$, that navigates a workspace $\mathcal{W} \subset \mathbb{R}^2$ for the purpose of classifying multiple fixed targets based on new (posterior) and prior sensor measurements, and environmental information. Both \mathcal{A} and \mathcal{S} are assumed to be rigid polygons, and \mathcal{S} has a fixed position and orientation with respect to \mathcal{A} . Prior measurements may be available from airborne sensors, environmental maps, or from another robotic sensor, and are used to estimate the geometry and location of targets and obstacles in \mathcal{W} [21, 32]. The robotic sensor workspace \mathcal{W} is populated with n fixed obstacles $B = \{\mathcal{B}_1, \dots, \mathcal{B}_n\} \subset \mathcal{W} \subset \mathbb{R}^2$, and r fixed targets $T = \{\mathcal{T}_1, \dots, \mathcal{T}_r\} \subset \mathcal{W} \subset \mathbb{R}^2$, with $\mathcal{B}_i \cap \mathcal{T}_j = \emptyset$ for $\forall i \in I_B$ and $\forall j \in I_T$, where I_B and I_T are the index sets of B and T .

Let \mathcal{F}_A be a moving Cartesian frame embedded in \mathcal{A} . Then, every point of \mathcal{A} and every point of \mathcal{S} have a fixed position with respect to \mathcal{F}_A , and a configuration vector $q = [x \ y \ \theta]^T$ can be used to specify the position (x, y) and orientation θ of both \mathcal{A} and \mathcal{S} with respect to a fixed inertial frame \mathcal{F}_W , embedded in \mathcal{W} . Obstacles and targets are also assumed to be fixed and rigid in \mathcal{W} , such that every point of \mathcal{B}_i , for $\forall i \in I_B$, and every point of \mathcal{T}_j , $\forall j \in I_T$, have a fixed position with respect to \mathcal{F}_W . Let the configuration space \mathcal{C} denote the space of all possible robot configurations.

A C-obstacle is a subset of \mathcal{C} that causes collisions with at least one obstacle in B , i.e., $\mathcal{CB}_i \equiv \{q \in \mathcal{C} \mid \mathcal{A}(q) \cap \mathcal{B}_i \neq \emptyset\}$, where $\mathcal{A}(q)$ denotes the subset of \mathcal{W} occupied by the platform geometry \mathcal{A} when the robot is in the configuration q . The union of all C-obstacles obtained from B is referred to as the C-obstacle region. Thus, in searching for targets in \mathcal{W} , the robotic sensor is free to rotate and translate in the free configuration space, which is defined as the complement of the C-obstacle region \mathcal{CB} in \mathcal{C} , i.e., $\mathcal{C}_{\text{free}} = \mathcal{C} \setminus \mathcal{CB}$ [33]. The path of the robotic platform's centroid is defined as a continuous map $\tau : [0, 1] \rightarrow \mathcal{C}$, with $q_0 = \tau(0)$ and $q_f = \tau(1)$, where q_0 and q_f are the initial and goal configurations, respectively. Since \mathcal{S} is mounted on \mathcal{A} , the path τ determines the targets in \mathcal{W} that can be measured by the robotic sensor, while traveling from q_0 to q_f .

The purpose for deploying the robotic sensor in \mathcal{W} is to obtain additional measurements to properly classify a subset of targets in T . It is assumed that to each target $\mathcal{T}_i \in T$ there is associated a discrete and, possibly, random variable, y_i , with finite range $\mathcal{Y} = \{y_i^1, \dots, y_i^p\}$, representing its classification. Due to limited sensor measurements or to targets being buried underground, y_i is non-observable or *hidden*, and must be inferred from a set of measurements, $M_i = \{m_{i1}, \dots, m_{if}\}$. Every measurement $m_{i\ell} \in M_i$ also is a discrete and random variable, with finite range $\mathcal{M}_{i\ell} = \{m_{i\ell}^1, \dots, m_{i\ell}^{N_{i\ell}}\}$, where $m_{i\ell}^k$ denotes the k th value of $m_{i\ell}$. After the robotic sensor obtains the set of measurements M_i from \mathcal{T}_i , y_i can be inferred from an available sensor model, which is typically given by a known joint probability mass function, $P(y_i, M_i)$.

While the platform \mathcal{A} must avoid collisions with the obstacles in B , the sensor's field-of-view \mathcal{S} must intersect $\mathcal{T}_i \in T$ in order to obtain the measurements M_i and infer y_i . Since \mathcal{S} is mounted on \mathcal{A} , the platform motion must be planned in concert with the sensor measurements, and the path τ must allow the robot to simultaneously avoid obstacles and search for the most valuable targets. Let the measurement set of a robotic sensor along a path τ be defined as $M(\tau) = \{M_i \mid \mathcal{T}_i \cap \mathcal{S}(q) \neq \emptyset, \tau(s) = q, s \in [0, 1], i \in I_T\}$, where $\mathcal{S}(q)$ is the subset of \mathcal{W} occupied by \mathcal{S} at a configuration q , along τ . Then, the robotic sensor path τ between q_0 and q_f must achieve multiple, simultaneous objectives:

1. Avoid all obstacles in \mathcal{W}
2. Minimize the distance traveled from q_0 to q_f
3. Maximize the information value of $M(\tau)$

In order to meet objectives 2 and 3, the robotic sensor performance is defined by an additive reward function that represents the profit of the information obtained along the path τ :

$$R(\tau) = w_V V(\tau) - w_D D(\tau) \quad (1)$$

Where, $V(\tau)$ is the information value of $M(\tau)$, and $D(\tau)$ is the distance traveled along τ , as defined in Section 4.3. The user-defined constants w_V and w_D weigh the trade-off between the value of the measurements and the distance traveled, respectively. Then, a robotic sensor path that meets objectives 1–3 is obtained by solving the following problem:

Problem 1 (Geometric Sensor Path Planning Problem) Given a layout \mathcal{W} and a joint probability mass function, $P(y_i, M_i)$, $\forall i \in I_T$, find a path τ^* for a robotic sensor with

platform \mathcal{A} and field-of-view \mathcal{S} that connects q_0 to q_f in $\mathcal{C}_{\text{free}}$, and maximizes the profit of information defined in Eq. 1.

The next section reviews existing probabilistic roadmap methods, and Bayesian sensor models that can be used to compute the measurements' information value. In Section 4, a new information roadmap deployment method is developed to solve Problem 1. The methodology is implemented on a demining application in Section 5, and its effectiveness is demonstrated through numerical simulations in Section 6.

3 Background

3.1 Probabilistic Roadmap Methods

Probabilistic roadmap methods (PRMs) are a class of randomized motion planning algorithms that have recently received considerable attention because they are capable of handling problems with many degrees of freedom, and large workspaces with many obstacles, for which other motion planning methods are computationally infeasible [23–26]. So far, these methods have been applied to classical motion planning, that is, to plan the path of a robot with geometry \mathcal{A} that avoids collisions with the set of obstacles B in \mathcal{W} , and connects q_0 to q_f by a minimum distance. In Section 4, PRM is modified to solve a geometric sensor path planning problem (Problem 1), accounting for the sensor's field-of-view geometry \mathcal{S} , and the set of targets T in \mathcal{W} to find the robotic sensor path τ^* .

PRM planners sample the free configuration space $\mathcal{C}_{\text{free}}$ and construct a roadmap graph by means of a *learning phase*. Subsequently, a collision-free path from q_0 to q_f is determined from the roadmap by a so-called *query phase* [26]. The roadmap is an undirected graph $G = (L, A)$, comprised of a set of nodes or *milestones* $L = \{c_1, \dots, c_{N_m}\}$, and a set of arcs $A = \{(c_i, c_j) \mid c_i, c_j \in L\}$. Every milestone $c_i \in L$ represents a value of the robot configuration q sampled from $\mathcal{C}_{\text{free}}$ using a probability density function (PDF). Every arc $(c_i, c_j) \in A$ represents a simple local path, typically a straight line, that connects c_i and c_j in $\mathcal{C}_{\text{free}}$, and is determined by an extremely fast local planner. Assuming the roadmap properly represents the connectivity of $\mathcal{C}_{\text{free}}$, multiple queries can be used in the second stage to construct a path from q_0 to q_f by concatenating several feasible paths in G , in order of increasing distance from nodes q_0 and q_f , until they are successfully connected. A smoothing algorithm can then be applied to obtain a more natural and shorter path by means of geometric operations [34].

The main difficulties in constructing appropriate roadmaps for PRMs are providing coverage of wide-open regions of $\mathcal{C}_{\text{free}}$, and representing the connectivity of $\mathcal{C}_{\text{free}}$ even in the presence of narrow passages. The use of uniform and Gaussian distributions have been proposed in [26] and in [25], respectively, in order to generate milestones that would properly cover wide-open regions of $\mathcal{C}_{\text{free}}$. More recently, an hybrid strategy combining a bridge test with a Gaussian probability density function along narrow passages, and a uniform probability density function in wide-open regions of $\mathcal{C}_{\text{free}}$ was proposed in [23] in order to construct appropriate roadmaps for workspaces with both characteristics. The bridge test is based on the observation that a narrow passage in \mathcal{C} has at least one restricted direction between C-obstacles

that can be identified by a short straight-line segment or *bridge* with endpoints on two C-obstacles. Let b and b' be two random variables representing the two endpoints of a bridge in \mathcal{C} . Since bridges must connect C-obstacles, b is sampled from a uniform distribution $f(b)$ over \mathcal{CB} that is zero in $\mathcal{C}_{\text{free}}$, and b' is sampled from the conditional PDF,

$$f(b' | b) = \lambda_b(b')I(b')/Z_b \quad (2)$$

normalized by $Z_b \equiv \int_{\mathcal{C}} \lambda_b(b')I(b')db'$. λ_b is a multivariate Gaussian PDF that is radially symmetric, and has a standard deviation that depends on the width of the bridge length (Section 5.2). $I(b)$ is a binary function that equals 1 when $b \in \mathcal{CB}$, and 0 otherwise. Then, as shown in [23], milestones can be obtained with a higher frequency inside narrow passages by sampling the PDF,

$$\pi_G(q) = \int_{\mathcal{C}} f(b' | b) f(b) db = \int_{\mathcal{CB}} \lambda_b(2q - b)I(2q - b)/Z_b db \quad (3)$$

obtained by choosing the desired configuration q as the midpoint on the bridge from b to b' , such that $b' = 2q - b$. In order to also include milestones from wide-open regions of \mathcal{C} , a uniform probability density function π_U is defined over $\mathcal{C}_{\text{free}}$, and the hybrid sampling strategy,

$$\pi_H = v\pi_G + (1 - v)\pi_U \quad (4)$$

is used to cover both wide-open regions and narrow passages in $\mathcal{C}_{\text{free}}$. The user-defined parameter $0 \leq v \leq 1$ is chosen to emphasize either distributions in the mixture (Section 5.2).

In Section 4, a new PDF is presented for sampling milestones based on the geometry and information value of the targets in \mathcal{W} . A new hybrid sampling strategy (Section 4.2) is derived from prior measurements and environmental information, and from the probabilistic sensor model reviewed in the next section. Then, the profit of information defined in Eq. 1 is used in place of the classical distance metric [24] in the query phase (Section 4.3) to determine the path of maximum information profit in the roadmap.

3.2 Bayesian Network Approach to Sensor Modeling

A common approach for modeling the sensor measurement process is to utilize a joint probability mass function (PMF) of the relevant variables, which may include target classification and features, sensor measurements and parameters (or mode), and environmental conditions. The joint PMF of a particular sensor may be obtained by means of estimation algorithms (e.g., [27, 30]), or by learning algorithms using wavelets or mixtures of Gaussians [16, 35]. In this paper, we adopt the method presented in [36, 37], in which the PMF is learned from data and represented by a Bayesian network (BN) model. The advantages of BN models are that they can easily deal with many variables, they are accompanied by very efficient learning and inference algorithms, and they provide a convenient factorization of the joint PMF that can be used to simplify the computation of posterior PMFs required by information theoretic functions (Section 4.1).

A Bayesian network model is a pair comprised of a directed graph and of a set of conditional probability tables (CPTs) that together specify the multivariate joint PMF of a set of discrete and random variables known as the *universe* [38]. Every random variable in the BN universe is assumed to have a finite range, and is represented by a node in the graph. Arcs between the nodes represent conditional probability relationships between the variables. As shown in [36, 37], the BN model of a sensor measurement process is obtained by defining the universe as the set of all variables that influence the sensor measurements, such as, the set of sensor-mode parameters S , the set of environmental conditions E , the set of measurements M , the set of hidden target features F to be inferred from M , and the hidden target classification variable y , i.e., $U \equiv \{S, E, M, F, y\}$. Then, the BN arcs and CPTs are determined from a database of prior sensor data, using BN batch learning algorithms [39]. The database consists of several cases in which all variables in U are sampled by obtaining sensor measurements from several known targets, under known operating and environmental conditions [36].


After the BN model is determined, it specifies the joint PMF underlying the sensor measurements in terms of the recursive factorization,

$$P(U_i) \equiv P(S_i, E_i, M_i, F_i, y_i) = \prod_{u_j \in U_i} P(u_j \mid pa(u_j)) \tag{5}$$

$$= P(M_i \mid S_i, E_i, F_i)P(F_i \mid y_i)P(y_i)P(S_i)P(E_i) \tag{6}$$

which holds for any target \mathcal{T}_i , i.e., $\forall i \in I_T$. Where, $pa(u_j)$ denotes the set of parents of a node $u_j \in U_i$, and factors in Eq. 5 are the BN CPTs. Since the parents of a node u_j are all the nodes in the BN with an outgoing arc to u_j , the factorization in Eq. 5 reflects the BN graph structure, which is learned from data and, thus, depends on the sensor type. The factorization in Eq. 6, corresponding to the BN structure in Fig. 1, has been shown to apply to various sensor types [36, 37]. When sensor measurements are available from a target \mathcal{T}_i , a junction-tree BN inference algorithm [40] computes the posterior PMF $P(F_i, y_i \mid \mathcal{E}_i)$ from the BN CPTs in Eq. 6, and from the evidence \mathcal{E}_i , which includes the values of the measurements M_i , and the values of known operating (S_i) and environmental conditions (E_i). The unknown target features and classification are estimated as the values with the highest posterior probability, i.e., $(\hat{F}_i, \hat{y}_i) = \operatorname{argmax}_{F_i^j, y_i^l} P(F_i = F_i^j, y_i = y_i^l \mid \mathcal{E}_i)$. And, the posterior probability $P(\hat{F}_i, \hat{y}_i \mid \mathcal{E}_i)$ is known as the confidence level (CL). The BN models of multiple and heterogeneous sensors can be used in combination with the Dempster–Shafer (DS) rule of evidence combination to perform feature-level sensor fusion [36]. Also, in this paper, the factorization in Eq. 6 is used to compute the information value of a measurement set before the measurements become available, as shown in Section 4.1.

Fig. 1 Typical structure of a BN sensor model [36]



4 Methodology: Information Roadmap Deployment (IRD)

PRM algorithms have been shown very effective at planning a collision-free path for robots with many degrees of freedom for the purpose of moving from an initial to a final configuration in a workspace by a minimum distance [23–26]. Along a separate line of research, several information-driven sensor planning algorithms have been developed to plan a sensor measurement process based on the expected information value of the measurements. The information roadmap deployment (IRD) presented in this section combines these two lines of research by using an information-value function to generate a new hybrid sampling strategy that increases the milestone density near targets with high information value, while also covering wide-open regions and narrow passages. As a result, IRD presents several advantages over existing sensor path planning approaches, such as, coverage path-planning [2, 11], random [2], grid [16], and probabilistic search strategies [1, 16, 17, 21, 22]. By extending the concept of a roadmap to the sensor planning problem, IRD can account for the geometries of the targets and of a moving sensor's field-of-view, and can be applied to robotic sensors with a finite platform geometry that navigate an obstacle-populated workspace. Also, thanks to the information-value function presented in the next section, IRD accounts for the influence of operating and environmental conditions on the measurements that can be obtained along the sensor path, prior to determining the features or classification of the targets that are inferred by following it.

4.1 Measurement Information Value

A basic difficulty in sensor planning consists of assessing the value or utility of the sensor measurements prior to obtaining them from the targets. Several information-theoretic functions have been proposed for this purpose. Cross entropy was used in [28] to solve a multisensor-multitarget assignment problem, and in [29, 30] to manage agile sensors with Gaussian models for target detection and classification. Entropy and the Mahalanobis distance measure were used in [27] for sensor selection in ad-hoc sensor networks. Recently, the authors showed that using entropy reduction leads to improved target classification and feature inference from multiple, heterogeneous sensor fusion, when the sensor models are not necessarily Gaussian [41]. As shown in this subsection, the expected entropy reduction (EER) can be used to represent the value of information, V , in terms of the joint PMF $P(y_i, M_i)$ known from the sensor model (Section 3.2). Other information-theoretic functions can be similarly applied, as shown in Section 4.2, to construct information roadmaps for other sensor applications, e.g., [27, 28, 30].

Entropy reduction is formulated in terms of conditional entropy, which can be used to represent the uncertainty in a discrete and random variable Y , given the value of another discrete and random variable Z , based on their joint and conditional PMFs,

$$H(Y|Z) = - \sum_z \sum_y P(y, z) \log_2 P(y|z) \quad (7)$$

where, \sum_y denotes marginalization over the range of Y [42]. Although information entropy is not additive, it can be shown [41] that the entropy reduction

$\Delta \hat{H}(Y; Z_j | Z_i) \equiv H(Y | Z_i) - H(Y | Z_i, Z_j)$ is additive, and represents the reduction in uncertainty brought about by Z_j , given prior information or evidence about Z_i .

In Problem 1 sensor measurements from a target \mathcal{T}_i are sought to reduce the uncertainty in the classification variable y_i . Thus, entropy reduction can be used to represent the value of a new (posterior) set of measurements M_i , given an a-priori evidence set \mathcal{E}_i^0 , which may include known environmental conditions, as well as the measurements and mode of a previously-deployed sensor. The superscript $(\cdot)^0$ denotes a set of random variables whose values are known a priori. Since the actual entropy reduction, $\Delta \hat{H}(y_i; M_i | \mathcal{E}_i^0)$, cannot be determined prior to measuring M_i , the expected entropy reduction (EER), defined as,

$$\Delta H(y_i; M_i | \mathcal{E}_i^0) \equiv H(y_i | \mathcal{E}_i^0) - \sum_{\ell=1}^f \sum_{k=1}^{N_\ell} [H(y_i | m_{i\ell} = m_{i\ell}^k) P(m_{i\ell} = m_{i\ell}^k | \mathcal{E}_i^0)], \quad \forall i \tag{8}$$

is used to represent the expected reduction in uncertainty in y_i that would be brought about by M_i , given \mathcal{E}_i^0 .

The conditional entropy $H(y_i | \mathcal{E}_i^0)$ in Eq. 8 is computed from the definition in Eq. 7, and from the posterior PMF $P(y_i | \mathcal{E}_i^0)$ computed by a junction-tree BN inference algorithm [40]. The conditional entropy $H(y_i | m_{i\ell} = m_{i\ell}^k)$ is computed from Eq. 7, using the posterior PMF,

$$\begin{aligned} P(y_i | m_{i\ell}^k) &= \frac{P(y_i)P(m_{i\ell}^k | y_i)}{P(m_{i\ell}^k)} \\ &= \frac{P(y_i) \sum_j \sum_l P(m_{i\ell}^k | f_{ij} = f_{ij}^l) P(f_{ij} = f_{ij}^l | y_i)}{\sum_j \sum_l P(m_{i\ell}^k | f_{ij} = f_{ij}^l) \sum_j P(f_{ij} = f_{ij}^l | y_i = y_i^l) P(y_i = y_i^l)}, \quad \forall \ell, \end{aligned} \tag{9}$$

where \sum_j and \sum_l denote marginalization over the range of all features $f_{ij} \in F_i$, and \sum_j denotes marginalization over the range of y_i (i.e., \mathcal{Y}). Now, all of the probabilities in Eq. 9 are known from the BN CPTs $P(y_i)$, $P(F_i | y_i)$, and $P(M_i | S_i, E_i, F_i)$. $P(M_i | S_i, E_i, F_i)$ can be used to compute $P(M_i | F_i)$ either by marginalization or by using the evidence \mathcal{E}_i^0 when available. Equation 9 is derived using the simplification,

$$\begin{aligned} P(M_i | y_i) &= \sum_j \sum_l P(M_i | f_{ij}^l, y_i) P(f_{ij}^l, y_i) \\ &= \sum_j \sum_l P(M_i | f_{ij}^l) P(f_{ij}^l, y_i), \quad \forall i \end{aligned} \tag{10}$$

obtained by noting that y_i and M_i are d -separated given F_i , because of the serial connection between the respective nodes (Fig. 1). Finally, the last term in Eq. 8 is computed as $P(M_i | \mathcal{E}_i^0) = \sum_j \sum_l P(M_i | f_{ij} = f_{ij}^l) P(f_{ij} = f_{ij}^l | \mathcal{E}_i^0)$, where $P(F_i | \mathcal{E}_i^0)$ is obtained using a junction-tree BN inference algorithm [40].

In the presence of multiple targets, the information value of a set of measurements $M(\tau)$ along the path τ is the cumulative EER,

$$V(\tau) = V[M(\tau)] \equiv \sum_{M_i \in M(\tau)} \Delta H(y_i; M_i | \mathcal{E}_i^0) \quad (11)$$

where, y_i , M_i , and \mathcal{E}_i^0 are the classification variable, measurements, and a-priori evidence set corresponding to target \mathcal{T}_i , respectively, and $M(\tau)$ is defined in Section 2. Since the BN sensor model holds for any target in \mathcal{W} , every term in the summation in Eq. 11 can be computed from the corresponding a-priori evidence set \mathcal{E}_i^0 using Eqs. 7–10. In the next subsection a new sampling strategy is presented that is based on a probability density function obtained from the EER in Eq. 8. By combining this information-based sampling distribution with π_H in Eq. 4, a roadmap is constructed that captures the information value of the targets as well as the connectivity of $\mathcal{C}_{\text{free}}$.

4.2 Learning Phase: Analysis of the Sampling Distribution

As shown in the previous subsection, EER can be used to compute the expected information value of the measurements from a target \mathcal{T}_i prior to obtaining the measurements M_i , and prior to determining the target's features and classification F_i and y_i . In order for a sensor with a bounded field-of-view \mathcal{S} to obtain the measurements M_i , \mathcal{S} must intersect the i th target geometry \mathcal{T}_i . Therefore, the subsets of \mathcal{W} where the sensor can make target measurements can be defined as follows:

Definition 1 (C-Target) The target \mathcal{T}_i in \mathcal{W} maps in the robot's configuration space, $\mathcal{C}_{\text{free}}$, to the C-target region $\mathcal{CT}_i = \{q \in \mathcal{C}_{\text{free}} \mid \mathcal{S}(q) \cap \mathcal{T}_i \neq \emptyset\}$.

The union of all the C-targets in $\mathcal{C}_{\text{free}}$ is the C-target region \mathcal{CT} , and the union of all C-targets corresponding to a robot configuration q is the set of C-targets $\mathcal{CT}(q) = \{\mathcal{CT}_i \mid i \in I_T, \mathcal{S}(q) \cap \mathcal{T}_i \neq \emptyset\}$. The measurement set that can be obtained in q is the set of measurements corresponding to these C-targets, i.e., $M(q) = \{M_i \mid i \in I_T, \mathcal{CT}_i \in \mathcal{CT}(q)\}$. Therefore, the information value of the measurements that can be obtained in a robotic sensor configuration $q \in \mathcal{C}_{\text{free}}$ is the cumulative EER,

$$V(q) = V[M(q)] \equiv \sum_{M_i \in M(q)} \Delta H(y_i; M_i | \mathcal{E}_i^0) \quad (12)$$

It follows that a high density of milestones with high information value can be obtained by sampling the PDF defined as,

$$\pi_V(q) = \frac{V(q)}{\int_{\mathcal{CT}} V(q) dq} \quad (13)$$

which is proportional to the EER function $V(q)$, and is normalized by the total EER of the C-target region.

The above PDF is combined with a uniform PDF π_U and a Gaussian PDF π_G obtained by the bridge test in order to construct a roadmap that also captures the connectivity of $\mathcal{C}_{\text{free}}$ by covering narrow passages and wide-open regions. The new hybrid strategy,

$$\begin{aligned} \pi &= v_2\pi_V + (1 - v_2)\pi_H \\ &= v_2\pi_V + v_1(1 - v_2)\pi_G + (1 - v_1)(1 - v_2)\pi_U \end{aligned} \tag{14}$$

is obtained from the sampling strategy in Eq. 4 and the PDF in Eq. 13. Where, $0 \leq v_1 \leq 1$ and $0 \leq v_2 \leq 1$ are user-defined parameters that are chosen to emphasize the narrow passages versus wide-open areas, and to emphasize information value versus connectivity, respectively. In the actual implementation, the PDF π is discretized into a probability mass function (PMF) over $\mathcal{W} \in \mathbb{R}^2$, as shown in Algorithm 1. The sampling distribution π obtained by Algorithm 1 is plotted in Fig. 2 for a simple example involving a robotic sensor that can translate freely but cannot rotate, and has a square platform geometry \mathcal{A} (grey), and a square field-of-view \mathcal{S} (white). This robotic sensor navigates the workspace shown in Fig. 2a, where the obstacles' geometry are shown in black, and the target geometries are shown in color patterns that are representative of their information value. The sampling distribution over \mathcal{C} is plotted in Fig. 2b. In this simple example, \mathcal{C} is equal to \mathcal{W} because $q = [x \ y]^T$. As plotted in Fig. 2b, C-obstacles (black) have zero sampling probability, and C-targets have sampling probability proportional to their information value (EER). As shown in Sections 5 and 6, the same approach is applicable to robotic sensors with more degrees of freedom, such as robots that are capable of rotating, as well as to any other geometries for \mathcal{W} , \mathcal{A} , and \mathcal{S} .

After the sampling distribution π is computed by Algorithm 1, it is sampled N_m times to obtain the set of milestones L that constitute the nodes of the roadmap $G = (L, A)$. Subsequently, the set of arcs A is obtained by a local planner that connects every milestone $c_i \in L$ to every other milestones with a straight line segment. The lengths of these straight-line segments are used to sort the milestones, such that only the k nearest milestones to c_i are elected as its candidate neighbors. Where, k is a user-defined parameter that is chosen based on N_m and on the complexity of \mathcal{W} .

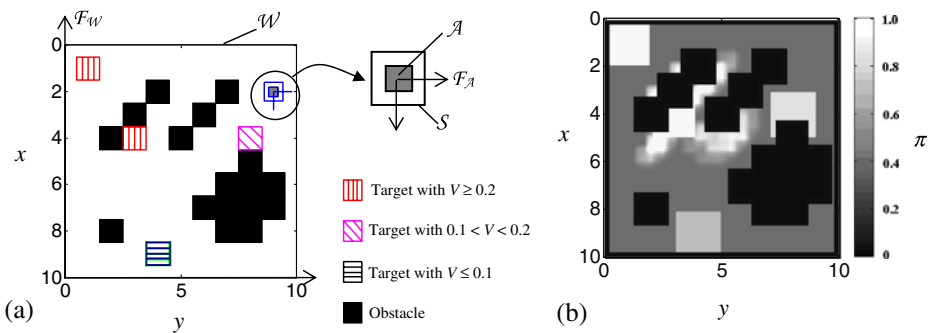


Fig. 2 Example of (normalized) sampling probability distribution π (b) for the non-rotating robotic sensor with platform \mathcal{A} and field-of-view \mathcal{S} , navigating the workspace \mathcal{W} in (a)

Algorithm 1 Sampling Distribution Generation

given $\mathcal{W}, \mathcal{A}, \mathcal{S}, P(y_i, M_i)$, and \mathcal{E}_i^0 for $\forall i \in I_T$
 discretize \mathcal{C} into N_q configurations $C = \{q^1, \dots, q^{N_q}\}$
 create an empty set of sample configurations, $Q = \emptyset$
 initialize variables, and let $I(q) = 0, U_Q = 0$ and $V_Q = 0$
for $\forall q^i \in C$ **do**
 if $q^i \in \mathcal{C}_{\text{free}}$ **then**
 $U(q^i) = 1$
 $I(q^i) = 0$
 compute $V(q^i)$ from Eq. 12
 else
 $U(q^i) = 0$
 $V(q^i) = 0$
 $I(q^i) = 1$
 end if
 for $\forall q_0^i \in C$ **do**
 compute $f(b' = q^i | b = q_0^i)$ from Eq. 2
 end for
end for
 $U_Q = \sum_{q^i \in C} U(q^i)$
 $V_Q = \sum_{q^i \in C} V(q^i)$
for $\forall q^i \in C$ **do**
 $\pi_U(q^i) = \frac{U(q^i)}{U_Q}$
 $\pi_V(q^i) = \frac{V(q^i)}{V_Q}$
 $\pi_G(q^i) = \sum_{q_0^i \in C} f(2q^i - q_0^i | q_0^i) I(2q^i - q_0^i) I(q_0^i)$
 $\pi(q^i) = v_2 \pi_V(q^i) + v_1(1 - v_2) \pi_G(q^i) + (1 - v_1)(1 - v_2) \pi_U(q^i)$
end for

For each candidate neighbor of c_i , the collision-check algorithm presented in [26] is implemented to check whether a simple path between them (typically a straight line segment) is collision free, by discretizing the path into a sequence of configurations. Then, the candidate neighbors with collision-free paths are connected to c_i to construct G . After G is constructed, it can be used to find the path with maximum information profit in Eq. 1 between any pair of q_0 and q_f , as explained in the next subsection.

4.3 Information-Driven Query Phase

After the information roadmap G is constructed by the method presented in the previous section, the query phase seeks a path τ^* from q_0 to q_f of maximum information profit (Problem 1). A query consists of connecting q_0 and q_f to G by searching for two milestones $c_0, c_f \in G$ that have the shortest collision-free distance to q_0 and q_f , respectively, and then finding the path from c_0 to c_f of maximum information profit. A path from c_0 to c_f in G is a sequence of adjacent cells that are pairwise connected by a simple collision-free path represented by the arc between

them, e.g., $\tau = \{c_0, c_i, c_j, \dots, c_f\}$, with $(c_i, c_j) \in A \subset G$, for any consecutive pair of indices in the index set I_τ of τ . Since every milestone represents a robot configuration, an additive distance metric can be defined for every pair $(c_i, c_j) \in A$, and used to compute the total path distance $D(\tau)$. A scaled Euclidian distance is adopted from [24] that changes the relative importance of position and orientation components of q through a set of weights organized in a diagonal and positive-definite matrix $W \in \mathbb{R}^{3 \times 3}$. Then, the total path distance is given by the sum of all weighted Euclidian norms along τ , i.e.,

$$D(\tau) = \sum_{i=0}^{f-1} \|W[q(c_{i+1}) - q(c_i)]\| \tag{15}$$

where, $q(c_i)$ denotes the robot configuration represented by c_i , and $\|\cdot\|$ represents the Euclidian norm [43, pp. 28–29].

Planners based on the A^* algorithm are the most effective at searching for the path of minimum total distance in G [24]. The A^* algorithm explores G iteratively, starting at c_0 and visiting every neighbor node c_i , to which a cost function is assigned by estimating the minimum-cost path from c_0 to c_f , through c_i [44, 45]. Based on the principle of optimality [46], this cost can be estimated as the sum of the actual cost of a path from c_0 to c_i , $g(c_i)$, plus an estimate of the minimum-cost path from c_i to c_f obtained from a heuristic function, $h(c_i)$. If the heuristic function is chosen such that $h(c_i)$ is always less than the actual cost of the (same) path from c_i to c_f , then the A^* algorithm is guaranteed to return a path of minimum cost whenever such a path exists in G [44, 45]. After a node is visited, the algorithm stores only the path of minimum cost, and labels the node as visited, assigning it a pointer to its parent node. This process forms a spanning tree T_s of the subset of G that has already been explored, and brings about considerable computational savings compared to other graph-searching algorithms [33, 47].

Since the goal of the robotic sensor is to maximize the measurement information profit, after connecting q_0 and q_f to G the information-driven query phase seeks the path τ^* with the maximum value of $R(\tau)$, defined in Eq. 1 as the trade-off of $V(\tau)$, defined in Eq. 11, and $D(\tau)$, defined in Eq. 15. Then, the A^* algorithm can be applied by defining the actual cost of a path $\tau_{0,i}^s$ connecting c_0 to c_i in the current spanning tree T_s as

$$g(c_i) \equiv -R(\tau_{0,i}^s) = -w_V \cdot V(\tau_{0,i}^s) + w_D \cdot D(\tau_{0,i}^s). \tag{16}$$

In order to guarantee that the cost estimated by the heuristic function is less than the actual cost, $h(c_i)$ is based on the cumulative EER of all measurements remaining after $\tau_{0,i}^s$, $\bar{M}(\tau_{0,i}^s)$, and on the length of a straight line segment $s_{i,f}$ connecting $q(c_i)$ to $q(c_f)$ in \mathcal{C} :

$$h(c_i) = - \sum_{\bar{M}(\tau_{0,i}^s)} w_V \cdot \Delta H(y_i; M_i | \mathcal{E}_i^0) + w_D \cdot D(s_{i,f}) \tag{17}$$

Where, the complement measurement set of a path is defined as $\bar{M}(\tau_{0,i}^s) = \{M_i | i \in I_T, M_i \notin M(\tau_{0,i}^s)\}$. The estimated total cost for a path from c_0 to c_f through c_i is $f(c_i) = g(c_i) + h(c_i)$. Thus, it includes the cumulative EER of all targets in \mathcal{W} . The actual set of measurements $M(\tau_{i,f})$ obtained along a path $\tau_{i,f}$ from c_i to c_f in G is a

subset of $\bar{M}(\tau_{0,i}^s)$. Then, since the EER in Eq. 8 always is positive by the properties of mutual information [42], the inequality,

$$\sum_{M_i \in \bar{M}(\tau_{0,i}^s)} \Delta H(y_i; M_i | \mathcal{E}_i^0) > \sum_{M_i \in M(\tau_{i,f})} \Delta H(y_i; M_i | \mathcal{E}_i^0), \quad \text{for } M(\tau_{i,f}) \subset \bar{M}(\tau_{0,i}^s), \quad (18)$$

holds for any path $\tau_{i,f}$ that is concatenated with $\tau_{0,i}^s$ to connect c_0 to c_f through c_i . Since the length of $s_{i,f}$ is always less than or equal to the shortest path between c_i and c_f in G , it also follows that $h(c_i) < -R(\tau_{i,f})$. Therefore, the A^* -type search in Algorithm 2 is guaranteed to always return the path τ^* of minimum cost or maximum information profit in the information roadmap G .

Algorithm 2 Pseudocode of A^* -search for τ^* in G

```

procedure  $A^*(G, c_0, c_f, R, h)$ 
  initialize  $T_s$  and OPEN as empty
  install  $c_0$  into  $T_s$ 
  OPEN := ( $c_0$ , OPEN); label  $c_0$  as visited
  while OPEN  $\neq \emptyset$  do
    find  $c := \arg \min_c (f(c))$  in OPEN
    if  $c = c_f$  then
      exit loop
    end if
    for every neighbor  $c'$  of  $c$  in  $G$  do
      if  $c'$  is not visited then
        OPEN := ( $c'$ , OPEN); mark  $c'$  as visited, and label it with a pointer to  $c$ 
      else
        if  $g(c') > g(c) + R(c, c')$  then
          modify  $T_s$  by redirecting the pointer of  $c'$  toward  $c$ 
          remove  $c'$  from OPEN
          OPEN := ( $c'$ , OPEN)
        end if
      end if
    end for
  end while
  if OPEN  $\neq \emptyset$  then
    return the minimum-cost path by tracing the pointers in  $T_s$  from  $c_f$  to  $c_0$ 
  else
    return failure
  end if

```

5 Implementation of Information Roadmap Deployment for Robotic Demining

The information roadmap deployment (IRD) methodology presented in the previous section is demonstrated through a demining system application in which the path of a robot with an on-board ground-penetrating radar (GPR) is planned based on

prior infrared (IR) sensor measurements and environmental information available from a minefield $\mathcal{W} \subset \mathbb{R}^2$. The purpose for deploying a GPR robotic sensor in $\mathcal{W} \subset \mathbb{R}^2$ is to infer the classification of targets that are either mines or clutter objects, and are buried in heterogeneous soils, under non-uniform environmental conditions [48]. The classification variable of every target, y_i , is non-observable and has two mutually-exclusive possible values $\mathcal{Y} = \{y_i^1, y_i^2\}$. y_i can be inferred from a set of measurements $M_i = \{d_{m_i}, z_{m_i}, s_{m_i}\}$ that pertain to the target features, namely, depth (d_i), size (z_i), and shape (s_i). The measured target features are typically extracted from raw sensor measurements through signal processing techniques, and thus contain random measurement errors and errors caused by unfavorable environmental conditions [49]. The actual target features $F_i = \{d_i, z_i, s_i\}$, together with y_i , can be inferred from M_i and any prior evidence \mathcal{E}_i^0 using a GPR BN sensor model described in Section 5.1. It is assumed that an IR sensor previously deployed on an airborne platform, such as an unmanned air vehicle (UAV), is used to obtain cursory measurements that together with other geospatial data (e.g., topography, land cover, and satellite imagery) provide \mathcal{E}_i^0 for $\forall i \in I_T$, as well as estimated targets' and obstacles' geometries, T and B , in \mathcal{W} . The GPR robotic sensor, on the other hand, makes more accurate measurements on the ground, but can only visit a subset of the targets in \mathcal{W} due to energy and time limitations.

The simulation of a demining sensor system developed in [36] is used to generate a rectangular minefield of chosen dimensions, or workspace, that includes several buried mines, clutter objects, obstacles, and heterogenous environmental conditions. A two-dimensional grid is superimposed on the minefield dividing it into square bins that are assigned a squared unit distance. Soil characteristics, vegetation, and time-varying meteorological conditions are modeled according to [50, 51], as shown in Table 1. The simulation assigns a set of environmental conditions to each bin, either at random or at user-specified positions. The targets are comprised of anti-tank mines (ATM), anti-personnel mines (APM), unexploded ordnance (UXO), and clutter objects (CLUT) that are sampled and reproduced using the Ordata Database

Table 1 Simulated minefield, sensor, and target characteristics

Symbol	Nodes	Range
y	Target classification	{clutter, mine, empty bin}
S	GPR mode: m_{GPR}	{depth search, resolution search, anti ground-bounce-effect search}
E	IR mode: m_{IR}	{surface-mine search, shallow-buried-mine search}
	Soil moisture (%): s_t	{dry [0, 10], wet (10, 40], saturated (> 40)}
	Soil composition: s_c	{very-sandy, sandy, high-clay, clay, silt}
	Soil uniformity: s_u	{yes, no}
	Vegetation: g	{no-vegetation, sparse, dense}
	Weather: w	{clear, overcast, raining}
F	Depth (cm): d	{low (7–10 a.m. and 6–9 p.m.), medium (10–1 p.m.), high (1–6 p.m.)}
	Size (cm): z	{surface [0], shallow-buried (0, 12], buried (12, 60], deep-buried (> 60)}
	Shape: s	{small (2, 13], medium (13, 24], large (24, 40], extra-large (> 40)}
	Shape: s	{cylinder, box, sphere, long-slender, irregular}

[52], which contains over 5,000 explosive items and 3,000 metallic and plastic objects that resemble anti-personnel mines. Each target occupies one or more bins in the minefield depending on its size (z_i), and is characterized by a depth (d_i) and a shape (s_i) that may take any of the values shown in Table 1. Thus, $\mathcal{T}_i \subset \mathcal{W}$ represents the geometry of the set of bins from which prior IR measurements are obtained for the i th target detected in \mathcal{W} . The ground robot is simulated using the nonholonomic unicycle model in $\mathcal{F}_{\mathcal{W}}$ [53, 54], with a platform geometry $\mathcal{A} \subset \mathcal{W}$ specified by the user. On-board the ground robot is a GPR sensor with a field-of-view geometry $\mathcal{S} \subset \mathcal{W}$, also specified by the user, that moves with the robot in $\mathcal{F}_{\mathcal{W}}$ and remains fixed with respect to \mathcal{A} . In the simulation, as soon as \mathcal{S} intersects a bin containing a target, measurements are reproduced and deteriorated based on the target features, the sensor mode and working principles, and the environmental conditions in the bin [36, 55].

5.1 GPR and IR Bayesian Network Models

The BN models of a GPR sensor (Fig. 3a) and of the Agema Thermovision 900 IR [36] are implemented in this paper to compute the expected information value of a path, defined in Eq. 11, or a robot configuration, defined in Eq. 12, and to fuse the sensors’ measurements after they are obtained. These BN models are learned from a database of sensor measurements using the approach presented in [36] and reviewed in Section 3.2. Because they rely on different operating principles and may function in different modes, these demining sensors are more or less effective depending on the environmental conditions. IR sensors, for example, detect anomalies in infrared radiation that is either emitted by mines, soil, or vegetation. Based on the location of the sensor, the radiation data can be processed to build an image of an horizontal area and to estimate the depth of the object therein for depths up to 12 cm. The mode m_{IR} influences the measured target features and is uniquely determined by its height above the ground. Therefore, airborne IR sensors typically obtain only cursory measurements of size z and shape s for shallow-buried objects. Because they rely on temperature variations, their performance also is highly influenced by illumination (time of day) i , weather w , vegetation g , and soil properties $s_r, s_c,$ and s_u , with ranges described in Table 1.




Fig. 3 Example of BN model (a), adapted from [36], and geometric characteristics (b) of a GPR robotic sensor

GPR sensors emit radio waves that penetrate the ground and process their reflections at the boundaries of materials characterized by different refraction indexes. Images of underground vertical slices and of any objects buried within are obtained over the field-of-view \mathcal{S} by sensing discontinuities in electrical properties. The measured size z_m , shape s_m , and depth d_m of an underground object can be obtained from these images through signal processing techniques such as edge extraction [49]. The frequency of the radio wave and its bandwidth determine the search mode m_{GPR} . Since penetration depth increases at lower frequencies and image resolution improves at higher frequencies, the optimal GPR mode depends on the target features F , and on the environmental conditions E shown in Table 1. For example, very high frequencies may be required in the presence of ground discontinuities to overcome the so-called ground-bounce effect (GBE) [50]. By providing complementary information about the targets, the GPR measurements can significantly improve target classification through feature-level fusion with prior IR data [36, 50]. As explained in Section 3.2 and [36], after τ^* is planned and executed, the GPR and IR BN models are used to estimate F_i and y_i for each target intersected by \mathcal{S} along τ^* using the a-posteriori evidence set \mathcal{E}_i . Here \mathcal{E}_i is comprised of fused GPR and IR measurements, and of the operating and environmental conditions, m_{GPR_i} , m_{IR_i} , and E_i .

5.2 Information Roadmap Deployment of Mobile GPR Sensor

Since the GPR sensor has a bounded field-of-view \mathcal{S} and is installed on a ground robot, its measurements depend on the robot path in \mathcal{W} . At the same time, the ground robot has a finite geometry \mathcal{A} that must avoid natural or man-made obstacles (e.g., water bodies, trees, and buildings) in \mathcal{W} . An example of robotic sensor geometry is shown in Fig. 3b. As explained in Section 2, the position and orientation of every point in \mathcal{S} and \mathcal{A} can be specified using one configuration vector, q , containing the coordinates and orientation of $\mathcal{F}_{\mathcal{A}}$ with respect to $\mathcal{F}_{\mathcal{W}}$. Airborne IR measurements are processed to obtain a map of the geometries and locations of potential obstacles B and targets T in \mathcal{W} . Together with the environmental information and the GPR BN model in Fig. 3a, the IR measurements are also used to compute the sampling PDF in Eq. 13, and to generate a set of N_m milestones by means of the hybrid sampling strategy in Eq. 14. N_m is defined by the user based on the size and complexity of \mathcal{W} , \mathcal{CB} , and \mathcal{CT} . The number of configurations N_q used to discretize \mathcal{C} (Algorithm 1) depends on the size of \mathcal{W} , and on the ranges of the robot's degrees-of-freedom. In this paper, the range $\theta \in [0, 7\pi/4]$ is discretized in eight intervals, and the ranges of x and y are discretized according to the two-dimensional grid superimposed on the minefield (Section 5). For example, for the workspace in Fig. 7, $N_q = 56.68 \cdot 10^3$ and $N_m = 900$.

The multivariate Gaussian λ_b can be generated using the Matlab[®] function *mvnrnd* [56], as follows: $\lambda_b = \text{mvnrnd}(\mu, \sigma)$, where the mean is $\mu = (x, y)$, and σ is a diagonal covariance matrix whose elements are chosen based on the width of the narrow passages from the interval [0.5, 2]. The user-defined parameters v_1 and v_2 in the sampling strategy (Eq. 14) determine the relative frequency of milestones sampled from narrow passages and from high information-valued targets, respectively. For example, the values used for minefields with a high density of obstacles and narrow passages (e.g., Fig. 9a) are $v_1 = 0.1$ and $v_2 = 0.3$. Similarly, the weights w_V and w_D in the reward function (Eq. 1) represent a trade-off between the value of

the measurements and the distance traveled, which typically is to be minimized to conserve time and energy. Thus, they depend on the application domain and on the units of V and D . Values of w_V and w_D in the intervals $[0.5, 1]$ and $[0, 0.5]$, respectively, were investigated in this paper, ultimately selecting $w_V = 0.9$ and $w_D = 0.1$ for the demining application.

The weighting matrix W in Eq. 15 is chosen as $W = \text{diag}(0.9, 0.9, 0.1)$ to scale the Euclidian distance based on the relative importance of the translational versus the rotational distance which, in turn, depends on the complexity of \mathcal{CB} [24]. Together with computational requirements, \mathcal{CB} also determines the number of candidate neighbor nodes k considered by the local planner, which in this paper takes integer values in the interval $[5, 20]$. In most of the simulations shown in Section 6, $k = 6$. However, for high-density obstacles minefields (Fig. 9a) the best value is found to be $k = 15$. For these parameters and the large workspace in Fig. 7 the maximum running time of the IRD algorithms was $1.5 \cdot 10^4$ s for the learning phase, and 155 s for the query phase on an Intel T2060 1.6 GHz CPU computer with 1.00 GB of RAM.

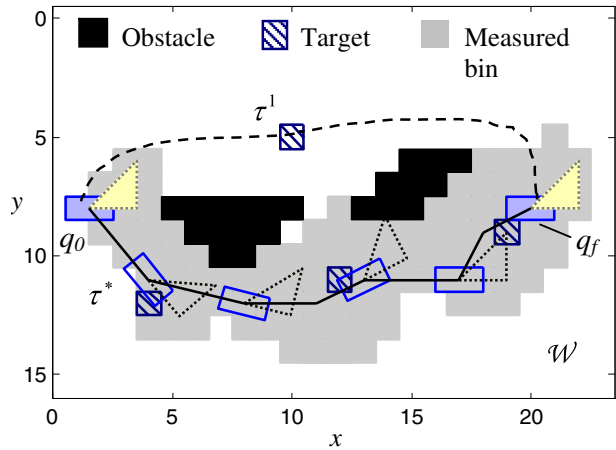
6 Results

The information roadmap deployment of a robotic GPR is tested on a variety of minefields exhibiting low-to-high densities of targets, obstacles, and narrow passages, as well as non-uniform soils, weather, and other environmental conditions. In Section 6.1, simple examples are used to illustrate and motivate the IRD approach, demonstrating the influence of the target geometries and prior information on the sensor path. In Section 6.2, the performance of IRD is compared to that of existing search strategies that are applicable to Problem 1 by means of the information roadmap, namely, shortest path [23–26], complete coverage [2, 11], and random [2] searches. Let $\Delta N_y(\tau)$ denote the number of targets that are correctly classified after fusing GPR and IR measurements along τ , minus the number of targets that were correctly classified based solely on IR measurements. Then the *classification efficiency* of a path τ is defined as the number of correctly-classified targets per distance traveled, $\eta_y(\tau) \equiv \Delta N_y(\tau)/D(\tau)$. After the path is executed and $M(\tau)$ becomes available, the *actual* entropy reduction, denoted by $\hat{V}(\tau)$, can be computed, and the function $\hat{V}(\tau)/D(\tau)$ can be used to assess the actual information value per distance traveled along a path τ . The results in Section 6.2 indicate that IRD outperforms existing approaches by up to one order of magnitude. Also, as shown in Section 6.2, IRD can be applied to plan the path of a non-overpass capable GPR platform, which must avoid collisions with mines as well as obstacles in order to prevent loss of the robotic sensor.

6.1 Influence of Prior Information and Workspace Geometry on the Sensor Path

In this section, a series of simple examples are used to demonstrate why the geometric characteristics of the problem, together with prior information, must be taken into account in planning the sensor path. In every example, the IRD sensor path τ^* is illustrated by plotting sample sensor configurations on the workspace. Other hypothetical paths are schematized by dashed lines for comparison. The first example

Fig. 4 Influence of target presence on τ^*



in Fig. 4 illustrates that both the location and geometry of targets and obstacles in \mathcal{W} must be accounted for in planning the path of a robotic sensor. Suppose the sensor must travel from q_0 to q_f , and \mathcal{W} contains two obstacles (black) and four equally-important targets (diagonal pattern) (Fig. 4). Although two obstacle-free paths τ^1 and τ^* of approximately the same distance can be found from q_0 to q_f , the path τ^1 (dashed line in Fig. 4) allows the GPR sensor to only make measurements from one target. Whereas, the IRD path τ^* allows the GPR sensor to make measurements from three of the targets in \mathcal{W} .

Besides accounting for targets’ geometries and locations, IRD also accounts for the expected information value of the measurements that can be obtained from them. Consider another simple example (Fig. 5) in which there exist three paths, τ^1 , τ^2 , and τ^* , that all allow the sensor to visit two targets by traveling approximately the same distance. Based on prior IR measurements, however, the information value of individual targets, defined in Eq. 8, is either medium (diagonal pattern) or low (horizontal lines), as shown in Fig. 5. The information value is discretized only for illustration purposes. If only the targets’ locations and geometries were taken into

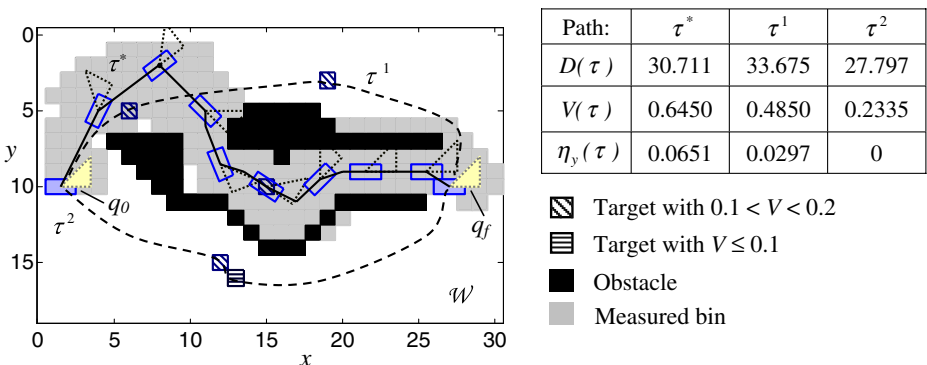


Fig. 5 Influence of prior sensor measurements on τ^*

account, these three paths would be considered equivalent. Instead, by maximizing the information profit, τ^* obtains a much higher classification efficiency than τ^1 or τ^2 .

The influence of the environmental conditions surrounding a target, E_i , on the sensor measurements is also accounted for by the value of information in Eq. 8. This is illustrated through an example in which the same targets buried in different soils lead to very different reduction of the uncertainty (entropy) in y_i , depending on how favorable the conditions are to the IR and GPR sensors. In the example in Fig. 6 two types of targets are buried in three different soil, for a total of six targets. In this workspace there are three paths from q_0 to q_f of approximately equal distance that each visit the same two targets buried in soils with different moisture. One target has high information value in saturated soil, and two targets have high information value in dry soil. As shown in Fig. 6, by visiting targets with high information value the GPR sensor obtains much higher classification efficiency along τ^* . Therefore, the environmental conditions influence the sensor path by making targets more or less valuable to the GPR, and IRD finds the path that enables the most valuable GPR measurements.

6.2 IRD Efficiency Comparison and Results

This section summarizes the results obtained by testing the IRD approach on a variety of minefields with various sizes, geometries, and environmental conditions. A representative example of GPR path computed by IRD is shown in Fig. 7 for a 64×108 (bin) minefield with 755 buried objects that include mines and clutter, concave polygonal obstacles, several narrow passages, and heterogeneous environmental conditions (not shown for simplicity). The bins measured by the GPR along τ^* (solid black line) are show in grey, and the resulting path efficiency is $\eta_y = 0.0913$. These results show that the robotic GPR is capable of navigating in an obstacle-populated workspace, and through narrow passages, in order to make measurements from targets with high information value (plotted in red and magenta in Fig. 7) with minimum distance.

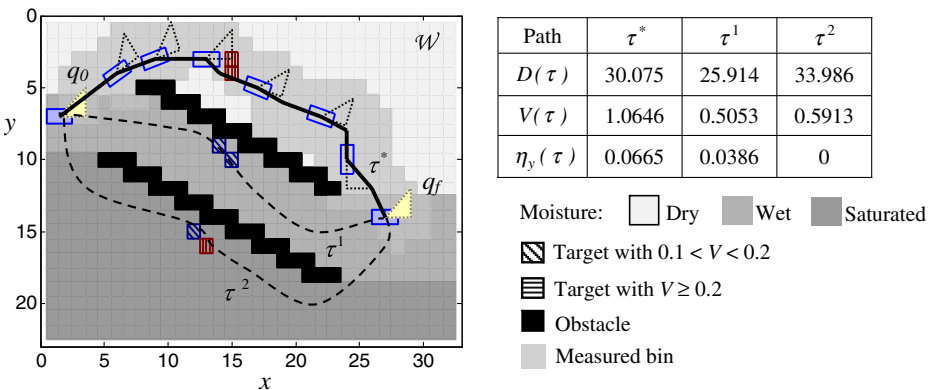


Fig. 6 Influence of environmental conditions on τ^*

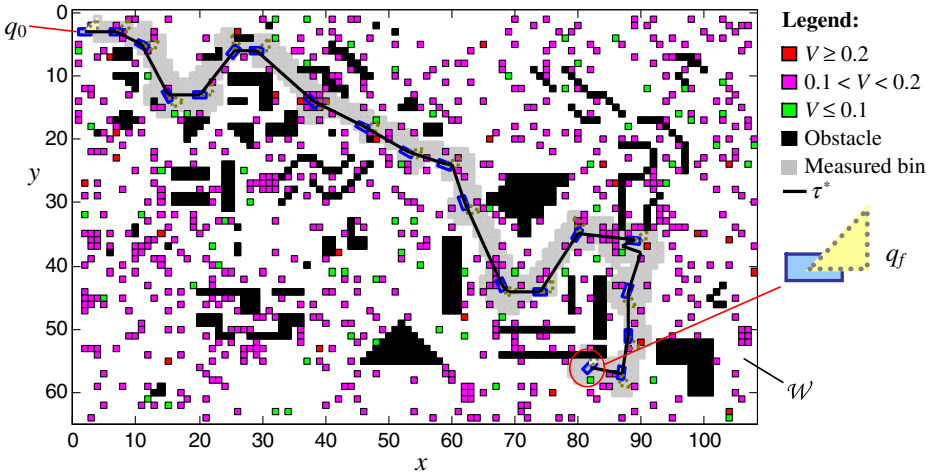


Fig. 7 Robotic GPR sensor path obtained by information roadmap deployment (IRD)

The average efficiency of IRD paths is compared to that of existing sensor path planning methods that are applicable to Problem 1, namely, shortest path [23–26], complete coverage [2, 11], and random [2] searches. Using the information roadmap presented in Section 4, these methods are implemented for the same workspace and robotic GPR sensor. The shortest path search [23–26] is a classical robot path planning strategy that is applied in the query phase (Section 4.3), and provides the value of the efficiency metric η_y when the presence of targets is not taken into account. By this approach, the roadmap G , obtained in Section 4.2, is searched for the path of minimum distance, τ_{short} , with an A^* -type algorithm that uses the straight-line distance as heuristic function, and does not consider the expected information value. Then, the GPR is turned on in a fixed mode along the path, and known environmental conditions are used as evidence in the GPR BN model to infer the targets’ classification from fused GPR-IR measurements.

Coverage path-planning algorithms play an important role in robotic sensing, because they emphasize the space swept by the robot’s sensor [11]. As pointed out in [11], one of the most significant sensor path planning results is a planner that generates a path that completely covers the obstacle-free space (e.g., see [1, 2, 22, 57]). For comparison, in this paper a complete-coverage path τ_{cover} is obtained for the robotic GPR in Fig. 3 by computing a lawn-mower-type path. This path consists of back-and-forth motions that avoid collisions between \mathcal{A} and the obstacles B , while making measurements from the entire free-configuration space \mathcal{C}_{free} with \mathcal{S} . Since computing the shortest complete-coverage path is known to be NP-hard [11], an ad-hoc solution to τ_{cover} is sought by placing milestones inside narrow passages, and near the boundaries of the C-obstacles and of the configuration space \mathcal{C} , denoted by $\partial\mathcal{C}B_i$ and $\partial\mathcal{C}$, respectively. The milestones inside narrow passages are obtained by the bridge test [23], and those near $\partial\mathcal{C}B_i$, $i \in I_T$, are obtained by a Gaussian sampler [25]. The milestones near $\partial\mathcal{C}$ are placed regularly spaced at a distance d_l that is based on the projection of \mathcal{S} onto $\partial\mathcal{C}$. The set of milestones is ordered by increasing x and y coordinates, as to reproduce a lawn-mower path in an

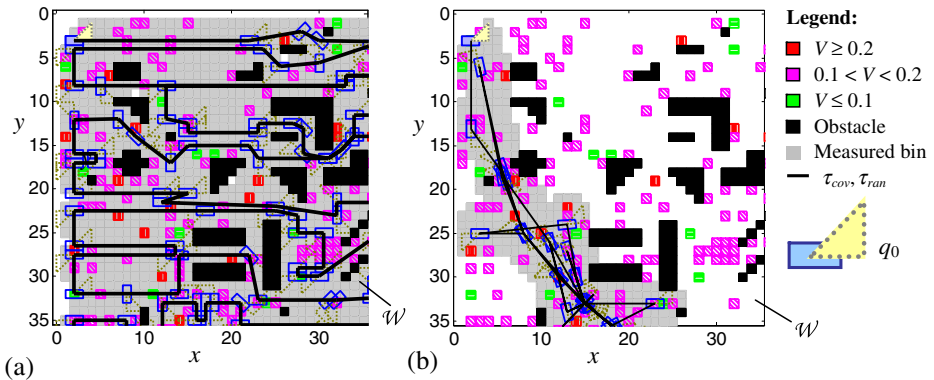


Fig. 8 Robotic GPR sensor paths obtained by complete coverage (a), and random search (b)

obstacle-free workspace. Then, the local planner described in Section 4.2 is used to connect the ordered milestones by collision-free straight-line segments. When the local planner fails (say in the x -direction), a milestone is inserted in the ordered list by increasing the other coordinate (say y) by a distance d_l . The new milestone is then connected to the nearest milestone in the list in the $-x$ direction, and so on. The result is a complete-coverage path that covers C_{free} , as illustrated by the grey bins measured by the GPR in Fig. 8b. The average efficiency of complete-coverage paths is shown in Table 2.

A popular approach in robotic demining and UXO clearance is to drive the robot in the minefield using a randomized search [2]. By this approach, a robot moves along a simple path, such as a straight line, until an obstacle is met, and then rotates a random amount before continuing along another simple path, while the on-board sensor is on at all times to detect targets along the path [2]. Using the roadmap G developed in Section 4.2, a random-search path, τ_{rand} , can be obtained by randomly selecting a node from the list of neighbors of c_0 in G , and then repeating the process for every new node until a pre-defined number of adjacent milestones connected by collision-free paths (arcs in G) is obtained. Although this approach does not optimize distance or information value in the query phase, it uses the information roadmap developed in Section 4.2 which contains a high density of configurations that enable high information-value measurements. Therefore, the resulting path τ_{rand} guarantees obstacle avoidance by \mathcal{A} , and \mathcal{S} is more likely to measure important targets than if deployed by a completely random strategy, such as the one presented in [2]. The four deployment strategies are applied to various minefields, for different values of q_0 and q_f , to obtain a representative average path efficiency η_y . The results are summarized

Table 2 Average sensor path efficiency

Metric	Deployment method			
	IRD	Shortest path	Random search	Complete coverage
$D(\tau)$	77	58	280	1,491
$\hat{V}(\tau)/D(\tau)$	0.119	0.0463	0.0243	0.0432
$\eta_y(\tau)$	0.2093	0.0175	0.0504	0.0693

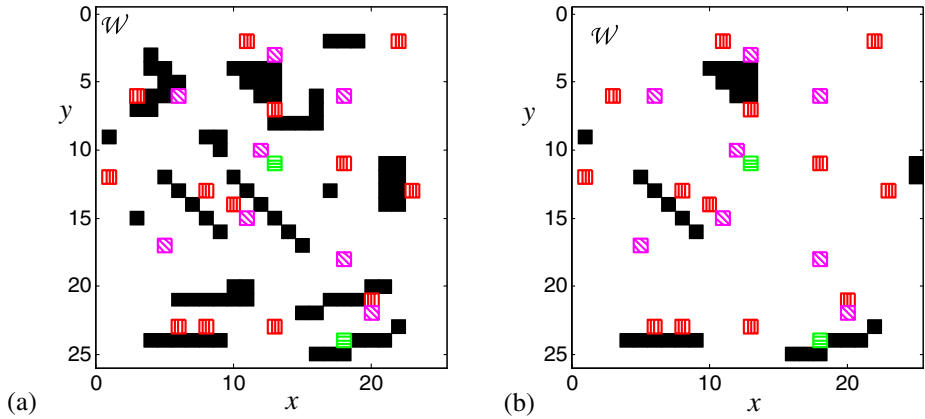
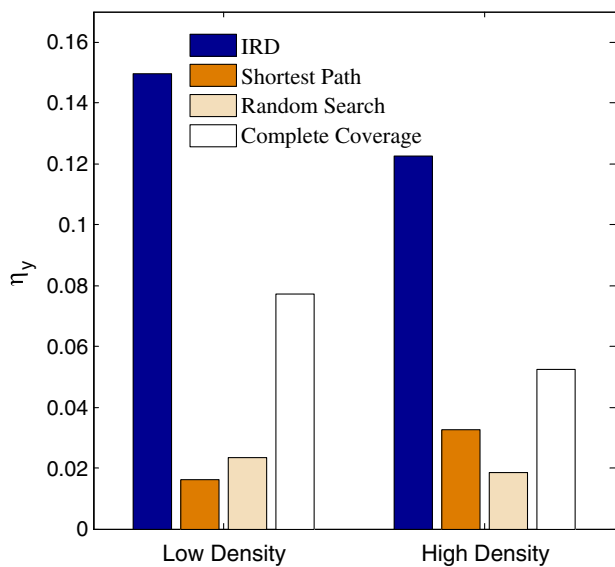


Fig. 9 Examples of workspaces with high **(a)** and low **(b)** obstacle density

in Table 2, and illustrate that IRD clearly outperforms the other methods leading to an average path efficiency three times greater than that of complete coverage, and one order-of-magnitude greater than that of shortest-path deployment. Additionally, these results show that the information roadmap developed in Section 4.2 also is very useful for implementing other search strategies, such as complete coverage, that may be valuable in other applications.

Since the sensor path efficiency depends on the characteristics of the workspace, extensive numerical simulations were performed to obtain average efficiency functions for different environmental conditions, and densities of obstacles and targets. Figure 9 illustrates two representative examples of workspaces that are considered

Fig. 10 Influence of obstacle density on GPR sensor path efficiency



to have high (a) or low (b) obstacle density. The target density is held constant in this study. As shown in Fig. 10, the path efficiency of IRD is significantly higher than that of shortest-path, random search, or complete coverage deployments, for both levels of obstacle density. The influence of environmental conditions on the method's efficiency is investigated by considering a minefield with the same geometric characteristics as the high obstacle-density example (Fig. 9a), but with two typologies of environmental conditions, such as the example shown in Fig. 11. The soil composition, moisture, and vegetation are plotted over the workspace in Fig. 11a for mild conditions, and in Fig. 11b for harsh conditions. As shown in Fig. 12, IRD achieves the best performance under both types of conditions, but its improvement compared to other methods is smaller for harsh environments, because the accuracy of the GPR measurements decreases regardless of the path. In fact, the shortest-path performance improves under harsh environmental conditions because when all

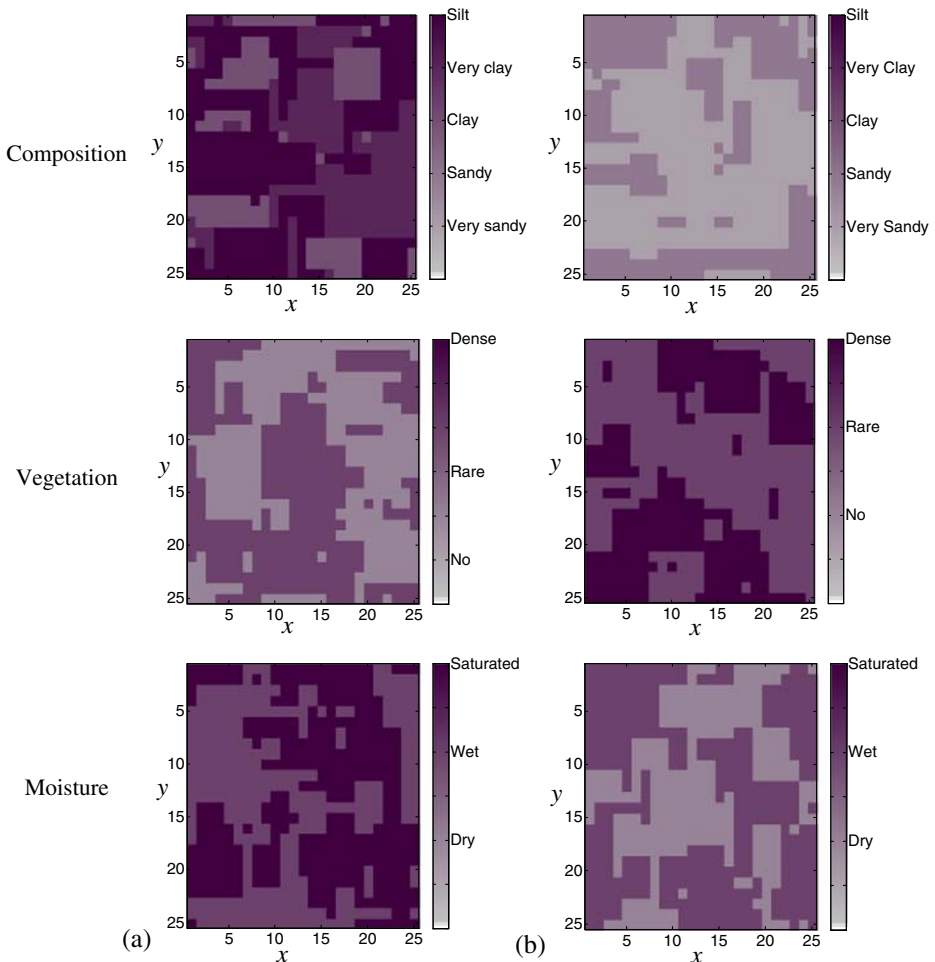
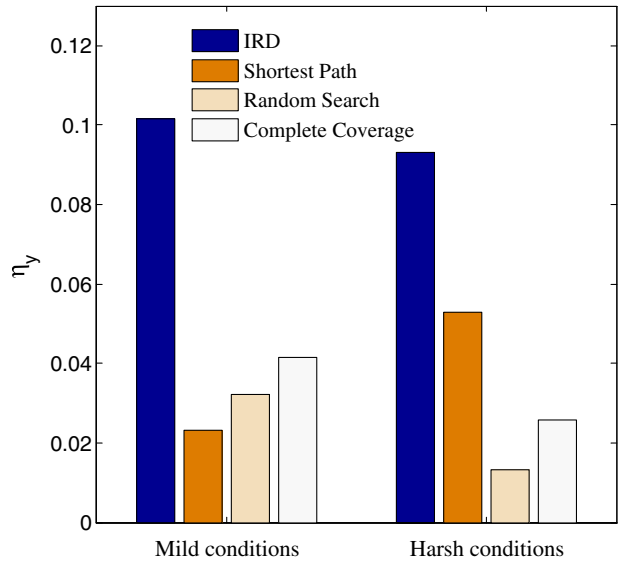


Fig. 11 Example of mild (a) and harsh (b) environmental conditions for \mathcal{W}

Fig. 12 Influence of environmental conditions on GPR sensor path efficiency



targets have low information value the efficiency metric is more heavily influenced by distance.

In another study, the average path efficiency is evaluated for low, medium, and high target densities (Fig. 13), using a constant obstacle density. As shown in Fig. 14, the IRD efficiency is significantly higher than that of complete coverage and random search methods, for all levels of target density. The most significant improvement is obtained in minefields with high target density because the sensor is able to visit more targets, taking full advantage of the proposed method. On the other hand, when the targets' density is low, their influence on the sensor path is considerably reduced, and therefore the efficiency of the IRD path approaches that of the shortest path. Conversely, the higher the target density, the closer the complete coverage efficiency will be to the IRD path, because the robotic sensor deployed by IRD will attempt to visit more targets within its reach.


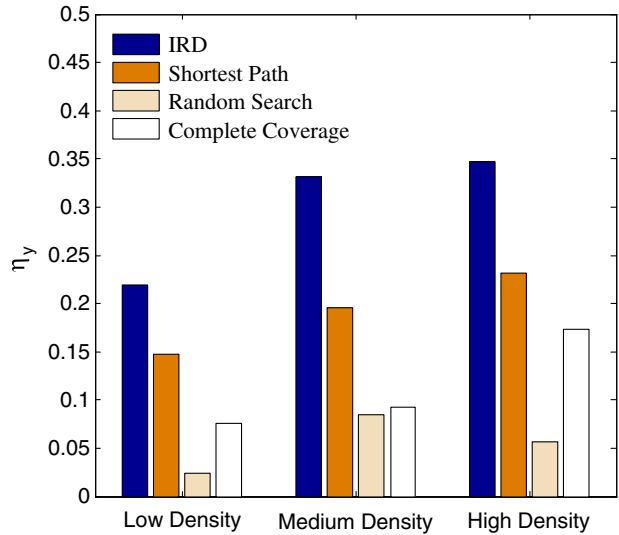


Fig. 13 Example of low (a), medium (b), and high (c) target density

Fig. 14 Influence of target density on GPR sensor path efficiency



Based on the results in Figs. 10–14 it can be concluded that the efficiency of IRD is significantly higher than that of existing robotic sensor planning approaches applicable to geometric sensing under a wide range of workspace conditions and characteristics. Another important application of the proposed method is planning the path of non-overpass capable robotic sensors that can be seriously damaged or even destroyed when driving over landmines [12]. For this type of robotic platforms, the set of targets T in \mathcal{W} must be treated as an additional set of obstacles, $B := \{B_1, \dots, B_n\} \cup T$, to be avoided by \mathcal{A} in case they are landmines. As shown by the examples in Fig. 15, the robotic GPR can still obtain measurements from the targets by navigating near them, such that they may be intersected by \mathcal{S} but avoided by \mathcal{A} . By implementing the information roadmap developed in Section 4.2, both IRD and

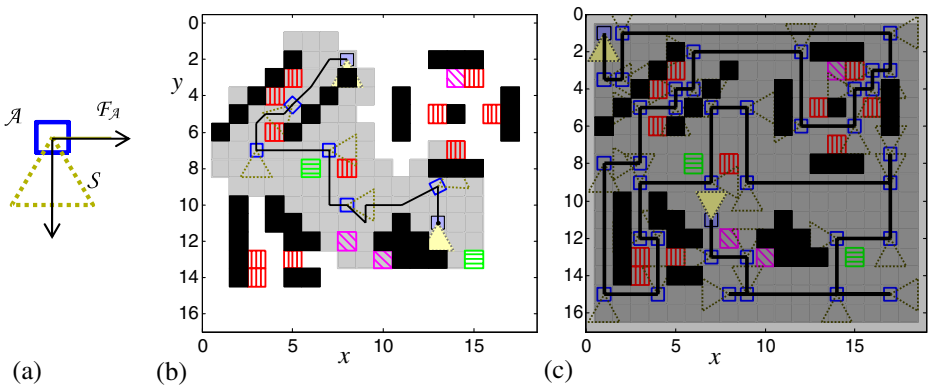
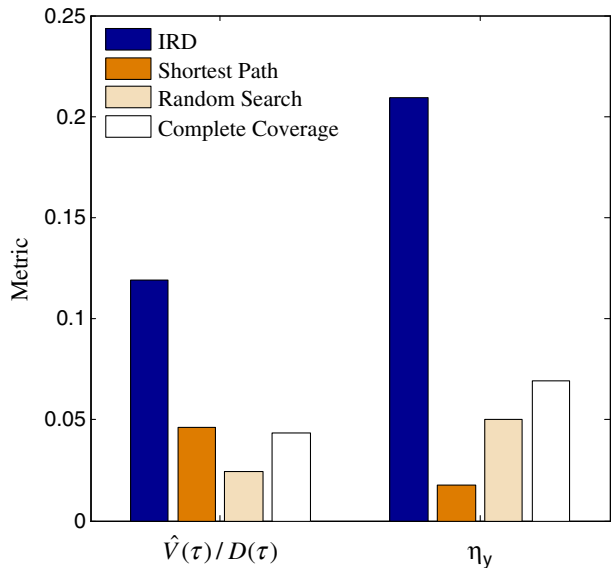


Fig. 15 Examples of IRD (b) and complete-coverage (c) paths for a non-overpass capable robotic platform \mathcal{A} , equipped with an on-board GPR sensor with field-of-view \mathcal{S} (a)

Fig. 16 Average sensor path efficiency for the non-overpass capable GPR sensor in Fig. 15a



complete coverage methods allow a non-overpass capable GPR sensor to navigate the workspace and make measurements from targets in \mathcal{W} , including those inside narrow passages which further restrict the free configuration space (Fig. 15). As can be expected, the average efficiencies (Fig. 16) are lower than those obtained by overpass-capable platforms (Table 2), because the sensor may need to travel a longer distance in order to avoid platform collisions with the targets. However, compared to shortest path, random search, and complete coverage, IRD still improves the number of targets that are properly classified per unit distance (η_y) by up to one order of magnitude (Fig. 16).

7 Summary and Conclusions

A novel information roadmap deployment (IRD) approach that combines information theory with probabilistic roadmap methods (PRMs) is presented and implemented in a simulated demining system. IRD computes a robotic sensor path that accounts for the geometry of the sensor's platform and field-of-view, and for the geometric characteristics of a workspace that is populated with multiple fixed targets and obstacles. The novel learning-phase and query-phase algorithms presented in Sections 4.2 and 4.3 use the targets' expected information value to generate a roadmap with a high density of high-information-value milestones, and a path that optimizes a desired trade-off between sensing performance and distance traveled. The value of information is quantified by the expected entropy reduction function presented in Section 4.1. By implementing an existing BN approach to model the sensor measurements (Section 3.2), the expected entropy reduction can be computed from the BN CPTs and from prior information, such as, prior sensor measurements and environmental conditions. Therefore, the couplings between environmental

and measurement uncertainty and the sensor's position are taken into account in planning the sensor path.

The IRD method is demonstrated by planning the path of a GPR sensor installed on a ground robot, based on prior measurements obtained by an airborne IR sensor, and environmental information, such as, soil characteristics, vegetation, and weather (Section 5). The results obtained from the simulations in Section 6 show that IRD outperforms existing sensor path planning methods applicable to geometric sensing, such as, complete coverage and random search, under a wide range of workspace conditions and geometric characteristics, increasing the average path efficiency by up to one order of magnitude. Also, the novel learning phase presented in Section 4.2 can be used to plan the path of non-overpass capable robotic sensors that can be seriously damaged or even destroyed when driving over landmines. As a result, IRD displays a classification efficiency several times greater than that of other methods, and the robotic sensor can make measurements even from targets located inside narrow passages which, in this case, further restrict the configuration space as they constitute potential obstacles.

Acknowledgements This work was supported by the Office of Naval Research Young Investigator Program (Code 321). We gratefully acknowledge Profs. Lawrence Carin and Ronald Parr at Duke University for their helpful guidance and suggestions.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Hofner, C., Schmidt, G.: Path planning and guidance techniques for an autonomous mobile cleaning robot. *Robot. Auton. Syst.* **14**, 199–212 (1995)
2. Acar, E.U.: Path planning for robotic demining: robust sensor-based coverage of unstructured environments and probabilistic methods. *Int. J. Rob. Res.* **22**, 441–466 (2003)
3. Kreucher, C., Kastella, K., Hero, A.: Multi-platform information-based sensor management. *Proc. SPIE* **5820**, 141–151 (2005)
4. Rao, N., Hareti, S., Shi, W., Iyengar, S.: Robot navigation in unknown terrains: introductory survey of non-heuristic algorithms. In: Technical Report ORNL/TM-12410. Oak Ridge National Laboratory, Oak Ridge, TN (1993)
5. Rao, N.: Robot navigation in unknown generalized polygonal terrains using vision sensors. *IEEE Trans. Syst. Man Cybern.* **25**(6), 947–962 (1995)
6. Lazanas, A., Latombe, J.C.: Motion planning with uncertainty—a landmark approach. *Artif. Intell.* **76**, 287–317 (1995)
7. Song, K., Chang, C.C.: Reactive navigation in dynamic environment using a multisensor predictor. *IEEE Trans. Syst. Man Cybern., Part B* **29**(6), 870–880 (1999)
8. Kazemi, M., Mehrandezh, M., Gupta, K.: Sensor-based robot path planning using harmonic function-based probabilistic roadmaps. In: Proc. ICAR '05, 12th International Conference on Advanced Robotics (2005)
9. Sun, Z., Reif, J.: On robotic optimal path planning in polygonal regions with pseudo-euclidian metrics. *IEEE Trans. Syst. Man Cybern., Part A* **37**(4), 925–936 (2007)
10. Lai, X.-C., Ge, S.-S., Al-Mamun, A.: Hierarchical incremental path planning and situation-dependent optimized dynamic motion planning considering accelerations. *IEEE Trans. Syst. Man Cybern., Part A* **37**(6), 1541–1554 (2007)
11. Choset, H.: Coverage for robotics: a survey of recent results. *Ann. Math. Artif. Intell.* **31**(1–4), 113–126 (2001)

12. Siegel, R.: Land mine detection. *IEEE Instrum. Meas. Mag.* **5**(4), 22–28 (2002)
13. Ferrari, S., Cai, C., Fierro, R., Perteet, B.: A multi-objective optimization approach to detecting and tracking dynamic targets in pursuit-evasion games. In: *Proc. of the 2007 American Control Conference*, pp. 5316–5321. New York, NY (2007)
14. Culler, D., Estrin, D., Srivastava, M.: Overview of sensor networks. *Computer* **37**(8), 41–49 (2004)
15. Juang, P., Oki, H., Wang, Y., Martonosi, M., Peh, L., Rubenstein, D.: Energy efficient computing for wildlife tracking: design tradeoffs and early experiences with zebnet. In: *Proc. 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)* (2002)
16. Liao, X., Carin, L.: Application of the theory of optimal experiments to adaptive electromagnetic-induction sensing of buried targets. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(8), 961–972 (2004)
17. Spletzer, J.R., Taylor, C.J.: Dynamic sensor planning and control for optimally tracking target. *Int. J. Robot. Res.* **22**(1), 7–20 (2003)
18. Hager, G.D., Mintz, M.: Computational methods for task-directed sensor data fusion and sensor planning. *Int. J. Robot. Res.* **10**, 285–313 (1991)
19. Chen, S.Y., Li, Y.F.: Automatic sensor placement for model-based robot vision. *IEEE Trans. Syst. Man Cybern., Part B* **34**(1), 393–408 (2004)
20. Chen, S.Y., Li, Y.F.: Vision sensor planning for 3cd model acquisition. *IEEE Trans. Syst. Man Cybern., Part B* **35**(5), 894–904 (2005)
21. Gelenbe, E., Cao, Y.: Autonomous search for mines. *Eur. J. Oper. Res.* **108**, 319–333 (1998)
22. Qian, M., Ferrari, S.: Probabilistic deployment for multiple sensor systems. In: *Proc. of the 12th SPIE Symposium on Smart Structures and Materials: Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, vol. 5765, pp. 85–96. San Diego (2005)
23. Sun, Z., Hsu, D., Jiang, T., Kurniawati, H., Reif, J.H.: Narrow passage sampling for probabilistic roadmap planning. *IEEE Trans. Robot.* **21**(6), 1105–1115 (2005)
24. Amato, N.M., Bayazit, O.B., Dale, L.K., Jones, C., Vallejo, D.: Choosing good distance metrics and local planners for probabilistic roadmap methods. *IEEE Trans. Robot. Autom.* **16**(4), 442–447 (2000)
25. Boor, V., Overmars, M.H., der Stappen, A.F.: The gaussian sampling strategy for probabilistic roadmap planners. *Proc. Conf. Robot. Autom.* **2**, 1018–1023 (1999)
26. Kavvaki, L.E., Svetska, P., Latombe, J.C., Overmars, M.H.: Probabilistic roadmaps for path planning in high-dimensional configuration space. *IEEE Trans. Robot. Autom.* **12**(4), 566–580 (1996)
27. Zhao, F., Shin, J., Reich, J.: Information-driven dynamic sensor collaboration. *IEEE Signal Process. Mag.* **19**, 61–72 (2002)
28. Schmaedeke, W.: Information based sensor management. In: *Proc. SPIE Signal Processing, Sensor Fusion, and Target Recognition II* (1993)
29. Kastella, K.: Discrimination gain to optimize detection and classification. *IEEE Trans. Syst. Man Cybern., Part A, Syst. Humans* **27**(1), 112–116 (1997)
30. Kreucher, C., Kastella, K., Hero, III, A.O.: Sensor management using an active sensing approach. *Signal Process.* **85**, 607–624 (2005)
31. Ji, S., Parr, R., Carin, L.: Nonmyopic multiaspect sensing with partially observable markov decision processes. *IEEE Trans. Signal Process.* **55**(1), 2720–2730 (2007)
32. Scrapper, C., Takeuchi, A.: Using a priori data for prediction and object recognition in an autonomous mobile vehicle. In: *Proc. SPIE Unmanned Ground Vehicle Technology*, vol. 5083 (2003)
33. Latombe, J.C.: *Robot Motion Planning*. Kluwer Academic, Boston (1991)
34. Berchtold, S., Glavina, B.: *A Scalable Optimizer for Automatically Generated Manipulator Motions*, pp. 1796–1802. München, Germany (1994)
35. Perrin, S., Duflos, E., Vanheege, P., Bibaut, A.: Multisensor fusion in the frame of evidence theory for landmines detection. *IEEE Trans. Syst. Man Cybern., Part C* **34**(4), 485–498 (2004)
36. Ferrari, S., Vaghi, A.: Demining sensor modeling and feature-level fusion by bayesian networks. *IEEE Sensors* **6**, 471–483 (2006)
37. Cai, C., Ferrari, S., Ming, Q.: Bayesian network modeling of acoustic sensor measurements. In: *Proc. IEEE Sensors*, pp. 345–348. Atlanta, GA (2007)
38. Jensen, F.V.: *Bayesian Networks and Decision Graphs*. Springer, New York (2001)
39. Heckerman, D., Geiger, D., Chickering, D.M.: Learning bayesian networks: the combination of knowledge and statistical data. *Mach. Learn.* **20**, 197–243 (1995)

40. Murphy, K.: How to use bayes net toolbox. [Online]. Available:<http://www.ai.mit.edu/~murphyk/Software/BNT/bnt.html> (2004)
41. Cai, C. Ferrari, S.: Comparison of information-theoretic objective functions for decision support in sensor systems. In: Proc. American Control Conference, pp. 63–133. New York, NY (2007)
42. Cover, T.M., Thomas, J.A.: Elements of Information Theory. Wiley, New York (1991)
43. Stengel, R.F.: Optimal Control and Estimation. Dover (1994)
44. Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cybern.* **4**(2), 100–107 (1968)
45. Nilsson, N.J.: Principles of Artificial Intelligence. Morgan Kaufmann, Los Altos, CA (1980)
46. Bellman, R.E., Dreyfus, S.E.: Applied Dynamic Programming. Princeton University Press, Princeton, NJ (1962)
47. Russell, S., Norvig, P.: Artificial Intelligence A Modern Approach. Prentice Hall, Upper Saddle River, NJ (2003)
48. Daniels, D.J.: A review of GPR for landmine detection. *Sens. Imaging* **7**(3), 90–123 (2006)
49. Paik, J.: Image processing-based mine detection techniques using multiple sensors: a review. *Subsurf. Sens. Technol. Appl.* **3**, 203–252 (2002)
50. MacDonald, J.: Alternatives for Landmine Detection. Rand, Sta Monica (2003)
51. Dam, R.V., Borchers, B., Hendrickx, J., Hong, S.: Soil effects on thermal signatures of buried nonmetallic landmines. In: Detection and Remediation Technologies for Mines and Minelike Targets VIII. Proc. of the SPIE, vol. 5089, pp. 1210–1218 (2003)
52. Explosive, Ordnance, Disposal, (EOD), and Technicians: ORDATA [Online]. Available:<http://maic.jmu.edu/ordata/mission.asp> (2006)
53. Fierro, R., Lewis, F.: Control of nonholonomic mobile robot using neural networks. *IEEE Trans. Neural Netw.* **9**(4), 589–600 (1998)
54. Cruz, D., McClintock, J., Perteet, B., Orqueda, O., Cao, Y., Fierro, R.: Decentralized cooperative control—a multivehicle platform for research in networked embedded systems. *IEEE Control Syst. Mag.* **27**, 58–78 (2007)
55. Vaghi, A.: Sensor management by a graphical model approach. Politecnico di Milano, Laurea Thesis (2004)
56. Mathworks: Matlab [Online]. Available:<http://www.mathworks.com> (2004)
57. Arkin, E.-M., Fekete, S.-P., Mitchell, J.-S.-B.: Approximation algorithms for lawn mowing and milling. *Comput. Geom.* **17**(1), 25–50 (2000)