

Random Disambiguation Paths for Traversing a Mapped Hazard Field

Carey E. Priebe,¹ Donniell E. Fishkind,¹ Lowell Abrams,² Christine D. Piatko³

¹ *Department of Applied Mathematics and Statistics, Whiting School of Engineering, Johns Hopkins University, Baltimore, Maryland 21218-2682*

² *Department of Applied Mathematics, George Washington University, Washington, DC 20052*

³ *Applied Physics Laboratory, Johns Hopkins University, Baltimore, Maryland 20723-6099*

Received 7 May 2003; revised 30 October 2003 and 10 September 2004; accepted 6 January 2005

DOI 10.1002/nav.20071

Published online 14 February 2005 in Wiley InterScience (www.interscience.wiley.com).

Abstract: We consider the problem of safely and swiftly navigating through a spatial arrangement of potential hazard detections in which each detection has associated with it a probability that the detection is indeed a true hazard. When in close proximity to a detection, we assume the ability—for a cost—to determine whether or not the hazard is real. Our approach to this problem involves a new object, the *random disambiguation path* (RDP), which is a curve-valued random variable parametrized by a binary tree with particular properties. We prove an admissibility result showing that there is positive probability that the use of an RDP reduces the expected traversal length compared to the conventional shortest zero-risk path, and we introduce a practically computable additive-constant approximation to the optimal RDP. The theoretical considerations are complemented by simulation and example. © 2005 Wiley Periodicals, Inc. *Naval Research Logistics* 52: 285–292, 2005.

Keywords: random path; path planning; marked point pattern; mine; minefield

1. INTRODUCTION

Suppose that a spatial point process has generated true and false detections, and each detection is marked with the probability that it is true. We wish to traverse a continuous curve of minimum arclength from a “source” point s to a “destination” point d (an “ s, d curve” of minimum arclength) avoiding “risk regions” about the true detections. Although the definitive true/false status of the detections is initially unknown, we assume the ability to “disambiguate” detections dynamically: When the curve is a specified distance from a detection, we have the option (for a fixed cost) of learning definitively whether the detection is false or true; we accordingly do or do not have the option to proceed through the associated risk region.

In Section 2 we introduce a new object, the *random disambiguation path* (RDP). This is an s, d curve-valued random variable parametrized by a *disambiguation protocol*—a rooted, binary tree whose vertices correspond to

disambiguation locations and whose branching covers all possible dynamically emerging results of the disambiguations. Our goal in this manuscript is to efficiently compute random disambiguation paths with small expected arclength, and to show their utility.

Shortest-path planning has been well-studied from the perspectives of graph theory, computational geometry, and robot motion planning (e.g., [4, 5, 7, 8]). In particular, path planning to avoid nonrandom disks in the plane can be accomplished efficiently using an associated visibility graph (discussed in Section 4.1). Our problem is more complicated than this classical problem, though, in that some of the disks may not be true hazards and may be entered when this status is confirmed. A related “expected path length” problem considered by Briggs, Scharstein, and Abbott [3] differs from ours in that it requires observing various landmarks before moving on to the next landmark. Their issues relate to obscuring of landmarks and waiting until a landmark becomes visible, and they use techniques related to a Markov Decision Process. The partially observed stochastic shortest path problem of Patek [11]—the extension of the stochastic shortest path problem of Bertsekas and Tsitsiklis [2] to the case of imperfect state information—is related to

Correspondence to: C.E. Priebe (cep@jhu.edu)

our RDPs; these authors approach their problem from a dynamic programming perspective.

Papadimitriou and Yannakakis have considered the shortest path problem when no prior map of obstacles is given [10]; they provide decision rules to optimize the ratio of the length of a path found dynamically, as obstacles such as unit squares are discovered, to the length of the statically optimal path. This contrasts with the situation considered here where the detection map is known in advance and each detection (obstacle) may turn out, with some associated probability, to not be an actual obstacle.

The structure of this manuscript is as follows: After introducing random disambiguation paths in Section 2, we prove an admissibility result in Section 3 showing that, under very mild assumptions, there is a positive probability with respect to the underlying point process that there will be an RDP strictly reducing expected traversal length compared to the conventional shortest zero-risk path. Computing the disambiguation protocol that minimizes the expected length of the associated random disambiguation path is computationally difficult, so in Section 4 we introduce a suboptimal RDP that is more practically computable, and we prove that it is within an additive constant of optimal. In Section 5 we perform simulations and explore a specific example. To simplify analysis and exposition we make several assumptions, including that the underlying space is \mathbb{R}^2 , that the risk regions are Euclidean balls of fixed radius about the detections, and that disambiguations are executed on the boundary of the corresponding risk regions. In Section 6 we briefly describe relaxations of these assumptions, as well as other possible directions for generalization.

2. RANDOM DISAMBIGUATION PATHS

Let \mathcal{S} be a marked point process on some simply connected, bounded subset $S \subseteq \mathbb{R}^2$. Unless otherwise specified, we will consider a particular realization of this process consisting of detections $x_1, x_2, \dots, x_n \in \mathbb{R}^2$, each x_i either being a *true detection* or a *false detection*. These observations are, respectively and independently, marked with $\rho_1, \rho_2, \dots, \rho_n \in (0, 1]$, where ρ_i is the probability that x_i is a true detection as rendered, for example, by the posterior probabilities of class membership from a Bayesian classifier [9, 12–15]. Let $[n] \equiv \{1, \dots, n\}$. Denote by $\mathcal{F} \subseteq [n]$ the set of indices of the true detections; the probability distribution of \mathcal{F} clearly follows from the independent marks ρ_i .

Let $\delta > 0$ be fixed and for $i = 1, 2, \dots, n$ denote by B_i the open ball in \mathbb{R}^2 of radius δ about x_i . For any deterministic $J \subseteq [n]$ and $x, y \in (\cup_{i \in J} B_i)^C$, the complement of the union of balls indexed by J , we denote by $q_{x,y,J}$ the x, y curve in $(\cup_{i \in J} B_i)^C$ of minimum arclength; modulo uniqueness issues (throughout this manuscript we assume the use of an arbitrary fixed selection rule to address ties and

alleviate uniqueness issues), $q_{x,y,J}$ is deterministic and efficiently computable using the notion of a visibility graph discussed in Section 4.1.

Let $s, d \in S^C$ be fixed. Our basic goal is to traverse a continuous s, d curve in $(\cup_{i \in \mathcal{F}} B_i)^C$ having arclength as small as we can practically attain. Because of the uncertainty of \mathcal{F} and the requirement that we avoid $\cup_{i \in \mathcal{F}} B_i$, it is not possible, without further knowledge of \mathcal{F} , to achieve an arclength less than that of $q_{s,d,[n]}$. However, in this manuscript we allow the possibility of dynamically *disambiguating* detections at a fixed cost $c \geq 0$ per disambiguation. Specifically, when a curve p originating at s is a distance of δ from some observation x_i —that is, the curve is on the boundary ∂B_i —we have the option of learning whether x_i is or is not a true detection and, accordingly, we may not or may proceed through B_i . When p terminates at d having avoided $\cup_{i \in \mathcal{F}} B_i$, the *traversal arclength* of p is defined to be $\ell(p) \equiv \ell^{\mathcal{E}}(p) + m \cdot c$, where $\ell^{\mathcal{E}}(p)$ is the Euclidean arclength of p and m is the number of disambiguations executed. This traversal arclength $\ell(p)$ is the objective quantity that we wish to minimize.

Let us first suppose that a maximum of one disambiguation may be executed. For any $j \in [n]$ and $z \in \partial B_j$ we introduce an s, d curve-valued random variable $p(z, j)$ called a *random disambiguation path* (RDP); this curve $p(z, j)$ first traverses the curve $q_{s,z,[n]}$, at which point x_j is disambiguated, and then traverses $q_{z,d,[n]}$ or $q_{z,d,[n] \setminus \{j\}}$, according as x_j is revealed to be a true or false detection. The expected length of $p(z, j)$ is

$$E\ell(p(z, j)) = \ell^{\mathcal{E}}(q_{s,z,[n]}) + \rho_j \ell^{\mathcal{E}}(q_{z,d,[n]}) + (1 - \rho_j) \ell^{\mathcal{E}}(q_{z,d,[n] \setminus \{j\}}) + c.$$

We also explicitly allow the choice $(z, j) = (d, \emptyset)$, corresponding to no disambiguations being executed, in which case $p(z, j) = p(d, \emptyset) = q_{s,d,[n]}$. The optimal RDP is $p^* \equiv p(z^*, j^*)$ where z^*, j^* satisfy

$$(z^*, j^*) \in \operatorname{argmin}_{(z,j)} E\ell(p(z, j)), \tag{1}$$

the minimization being over all pairs (z, j) described above.

In the more general case the maximum number of disambiguations allowed is a given positive integer K . A *disambiguation protocol* is a rooted binary tree T of depth K or less, in which each nonleaf¹ vertex v has both a left and right successor, and associated with each nonleaf vertex v is a pair $(z_v, j_v) \in \partial B_{j_v} \times [n]$. The *random disambiguation path* $p(T)$ is an s, d curve-valued random variable parametrized by T : Each root-leaf path v_1, v_2, \dots, v_s in T corresponds to the possible sequential disambiguations of detec-

¹ In the absence of other vertices we consider the root to be a leaf.

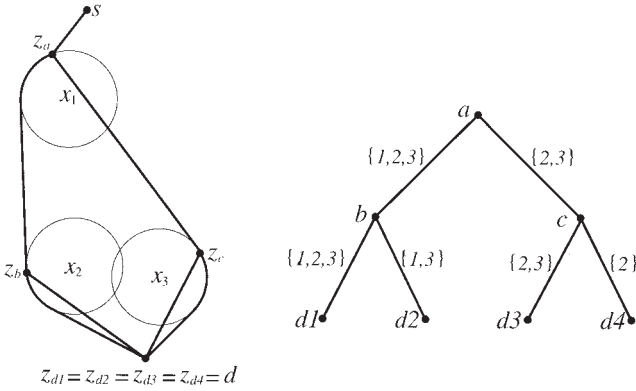


Figure 1. An RDP. Depicted on the left are three detections, their associated risk regions, and the various s, d curves which are the possible values of $p(T)$. On the right is the corresponding parametrizing tree. The edges are labeled with the appropriate set J_k indicating which detections are in fact true hazards or remain ambiguous.

tions indexed $j_{v_1}, j_{v_2}, \dots, j_{v_{s-1}}$ at respective locations $z_{v_1}, z_{v_2}, \dots, z_{v_{s-1}}$ with the outcome, for all k , that detection $x_{j_{v_k}}$ is true or false according as v_{k+1} is a left or right successor of v_k . From the k th disambiguation to the $(k + 1)$ st disambiguation $p(T)$ traverses $q_{z_{v_k}, z_{v_{k+1}}, J_k}$, where J_k is the set of indices of those detections not yet known to be false. (Here, we adopt the convention $z_{v_0} \equiv s$ and $z_{v_s} \equiv d$.) The dynamically realized sequential results of the disambiguations indicated by the protocol T determine a corresponding root-leaf path in T , which corresponds in turn to the s, d curve taken by $p(T)$. See Figure 1 for an example of an RDP.

The optimal RDP, denoted p^* , is $p(T^*)$, where

$$T^* \in \operatorname{argmin}_T \operatorname{El}(p(T)), \tag{2}$$

the minimization being over all possible protocols T described above. In general, p^* is not practically computable since the optimization problem in (2) is particularly difficult, even when $K = 1$. In Section 4 we show a way to compute a suboptimal RDP by restricting the domain of the optimization problem in (2) to a well-chosen finite set.

3. ADMISSIBILITY

This section presents an admissibility result stating that, under mild assumptions, an RDP provides, with positive probability, an expected strict improvement over the default $q_{s,d,[n]}$ which would be used if the disambiguation capability were not available.

We make two assumptions: (1) The marks ρ_i , conditioned on the associated observations x_i being true (false) detections, are independent and identically distributed with

a probability density function having support $(0, 1)$. (2) With respect to the underlying point process \mathcal{S} , the number of balls B_i that intersect the line segment $\overline{s, d}$ is almost surely finite, and with a positive probability at least one such intersection occurs.

THEOREM 1: Under the assumptions above on the spatial point process \mathcal{S} , there exists a disambiguation cost $c^* > 0$ and an allowed number of disambiguations $K > 0$ such that for all $c \in [0, c^*)$

$$P_{\mathcal{S}}[\operatorname{El}(p^*) < \ell^{\mathcal{E}}(q_{s,d,[n]})] > 0,$$

where $P_{\mathcal{S}}$ denotes probability relative to the underlying point process \mathcal{S} .

In other words, for a reasonable class of processes \mathcal{S} , the random disambiguation path p^* is, with positive probability, strictly superior to the conventional shortest zero-risk path $q_{s,d,[n]}$. [The value of the probability of superiority depends on the spatial point process \mathcal{S} ; processes can be specified for which this probability takes any value in $(0, 1)$.] Thus, path planners should consider the disambiguation option when the capability exists.

PROOF: Consider some $K > 0$ such that there is positive probability of exactly K balls B_i intersecting the line segment $\overline{s, d}$, and without loss of generality relabel the observations so that these intersecting balls are $\{B_i\}_{i=1}^K$. Consider the RDP p which sequentially disambiguates the detections it encounters along $\overline{s, d}$; if one of them is a true detection, then p retreats along $\overline{s, d}$ to s , then follows $q_{s,d,[n]}$ to d . We have

$$\begin{aligned} \operatorname{El}(p^*) \leq \operatorname{El}(p) &\leq \left(1 - \prod_{i=1}^K (1 - \rho_i)\right) (2\ell^{\mathcal{E}}(\overline{s, d}) + \ell^{\mathcal{E}}(q_{s,d,[n]})) \\ &\quad + \left(\prod_{i=1}^K (1 - \rho_i)\right) \ell^{\mathcal{E}}(\overline{s, d}) + Kc. \end{aligned}$$

Using a standard measure-theoretic argument, our assumptions imply that there is positive probability that there exists an $\epsilon > 0$ such that $\ell^{\mathcal{E}}(q_{s,d,[n]}) > \ell^{\mathcal{E}}(\overline{s, d}) + \epsilon$. Moreover, our assumptions imply that for all $\delta > 0$ there is a positive probability (with respect to \mathcal{S}) that $\prod_{i=1}^K (1 - \rho_i) > 1 - \delta$. Clearly, we have

$$\operatorname{El}(p^*) \leq \delta(2\ell^{\mathcal{E}}(\overline{s, d}) + \ell^{\mathcal{E}}(q_{s,d,[n]})) + \ell^{\mathcal{E}}(q_{s,d,[n]}) - \epsilon + Kc.$$

We may choose δ and c^* such that $\delta(2\ell^{\mathcal{E}}(\overline{s, d}) + \ell^{\mathcal{E}}(q_{s,d,[n]})) - \epsilon + Kc < 0$, from which the desired result follows. \square

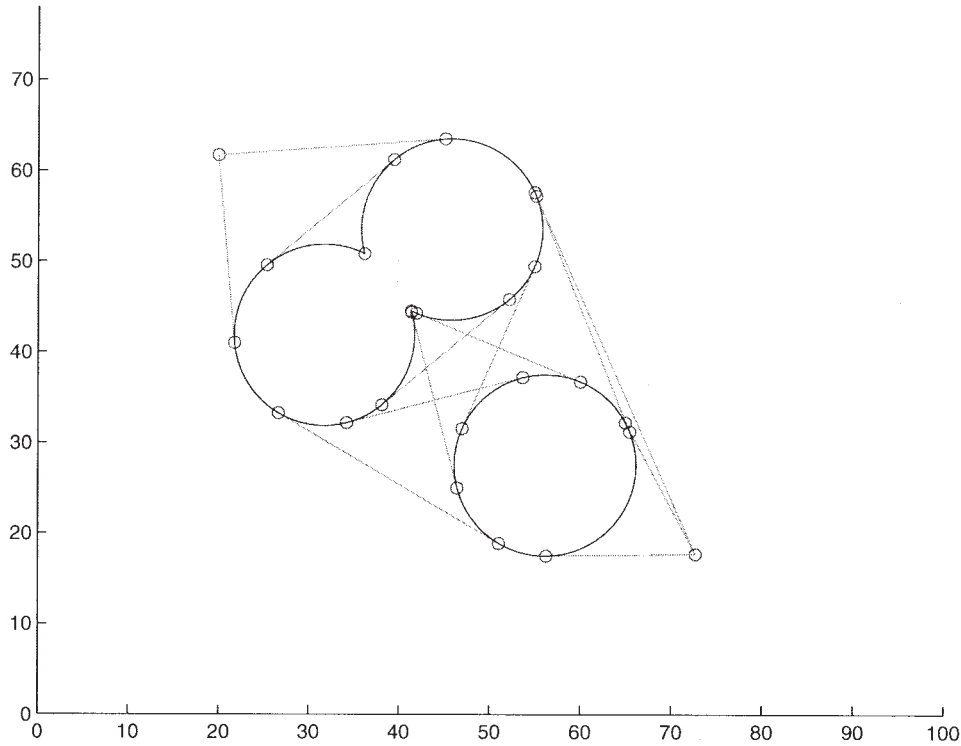


Figure 2. An example visibility graph, with vertices, arc edges, and line edges. Note that the two lightly-drawn (yellow) segments of disk boundary (in the intersection of disks) are *not* edges in the visibility graph.

4. APPROXIMATING THE OPTIMAL RDP

In general, p^* is not practical to compute due to the nature of the optimization problem in (2). In this section we consider restricting the domain of the optimization problem (2) to a finite and more practical subset, thus generating a suboptimal RDP that is within an additive constant of the optimal RDP p^* . We first define the *visibility graph*. This is used to compute $q_{x,y,J}$ and thus forms the core of the subroutine in any RDP algorithm for navigating from one disambiguation location to the next. The visibility graph will also be used to define our suboptimal RDP.

4.1. The Visibility Graph

Let $J \subseteq [n]$, $x, y \in (\cup_{i \in J} B_i)^C$ be specified. For distinct points $a, b \in \{x, y\} \cup \partial(\cup_{i \in J} B_i)$, the closed line segment $\overline{a, b}$ is a *tangent segment* provided that: (1) For all $r \in \{a, b\} \setminus \{x, y\}$, a, b is *tangential* to $\partial(\cup_{i \in J} B_i)$ at r , and (2) the relative interior of a, b is contained in the interior of $[(\cup_{i \in J} B_i) \cup \{x, y\}]^C$.

The *visibility graph* $G_{x,y,J} = (V_{x,y,J}, E_{x,y,J})$ is defined as follows. The vertex set $V_{x,y,J}$ of $G_{x,y,J}$ consists of: x, y , all points of $\partial(\cup_{i \in J} B_i)$ which intersect a tangent segment, and all points of $\partial(\cup_{i \in J} B_i)$ at which two or more ∂B_i 's intersect (for $i \in J$). The edge set $E_{x,y,J}$ of $G_{x,y,J}$ consists

of all tangent segments and all connected components of $\partial(\cup_{i \in J} B_i) \setminus V_{x,y,J}$ (these are circular arcs). The graph theoretic endpoints of these edges are their line and arc endpoints, respectively, and each edge is weighted with its arclength. See Figure 2 for an example visibility graph.

It is a well-known (and true) folk theorem that $q_{x,y,J}$ is the shortest x, y path in $G_{x,y,J}$. Since every pair of nonidentical ∂B_i 's have at most four mutually tangential lines and two points of intersection, the size of $V_{x,y,J}$ and $E_{x,y,J}$ are each $O(n^2)$. Thus the naive construction of $G_{x,y,J}$ requires $O(n^3)$ assignment, arithmetic, and trigonometric operations, and Dijkstra's algorithm with a heap implementation applied to $G_{x,y,J}$ yields $q_{x,y,J}$ in $O(n^2 \log n)$ operations (see, e.g., [1], pp. 115–116).

For each $j \in J$, the first point where $q_{x,y,J \setminus \{j\}}$ intersects ∂B_j (if this intersection exists) will be of interest; let $V'_{x,y,J}$ denote the union of $V_{x,y,J}$ and all such points (i.e., over all $j \in J$).

4.2. Approximating the Optimal RDP When $K = 1$

We begin with the case $K = 1$. Consider the minimization problem in Eq. (1) with its feasible region restricted to pairs $(z, j) \in (\partial B_j \cap V'_{s,d,[n]}) \times [n]$ [as above, we also explicitly include the choice $(z, j) = (d, \emptyset)$]. Let (z^*, j^*)

be the solution to this restricted problem, and let \hat{p}^* denote the RDP $p(\hat{z}^*, \hat{j}^*)$.

We can compute (\hat{z}^*, \hat{j}^*) and \hat{p}^* in $O(n^4 \log n)$ operations since there are $O(n^2)$ pairs (z, j) in the feasible region, each requiring the computation of the curves $q_{s,z,[n]}$, $q_{z,d,[n]}$, and $q_{z,d,[n]\setminus\{j\}}$ in time $O(n^2 \log n)$, and all of the needed visibility graphs can be constructed initially and simultaneously in time $O(n^4)$.

Recall that δ is the radius of the potential hazard detection regions and K is the allowed number of disambiguations. The next result shows that for $K = 1$ the suboptimal \hat{p}^* is within an additive constant, equal to the circumference of one of the hazard regions, of optimal.

LEMMA 2: When $K = 1$, we have $El(\hat{p}^*) \leq El(p^*) + 2\pi\delta$.

PROOF: By our construction of $G_{s,d,[n]}$, the distance in $(\cup_{i=1}^n B_i)^C$ from any point in $\partial \cup_{i=1}^n B_i$ to its nearest vertex in $V_{s,d,[n]}$ is not more than $\pi\delta =$ half the circumference of the B_i 's. In particular, if (z^*, j^*) is as in (1), and $(z^*, j^*) \neq (d, \emptyset)$, then there is some $z \in \partial B_{j^*} \cap V_{s,d,[n]}$ not more than $\pi\delta$ distant from z^* . The RDP $p(z, j^*)$ —regardless of \mathcal{F} —is thus not more than $\pi\delta + \pi\delta$ longer than $p(z^*, j^*)$; one $\pi\delta$ for the s, z part of the s, d curve and one $\pi\delta$ for the z, d part of the s, d curve. Now, $El(p(\hat{z}^*, \hat{j}^*)) \leq El(p(z, j^*)) \leq El(p(z^*, j^*)) + 2\pi\delta$, and the result is shown. \square

4.3. Approximating the Optimal RDP When $K > 1$

When $K > 1$, we can generalize the arguments in the preceding subsection. We say a protocol T is a *visibility protocol* if it meets the criterion that every nonleaf vertex v is assigned a pair $(z_v, j_v) \in (\partial B_{j_v} \cap V'_{z_w,d,J}) \times [n]$, where w is the predecessor of v (if v is the root, then z_w is defined to be s), J is the set of indices of detections not known to be false detections immediately prior to the disambiguation associated with v , and $j_v \neq j_u$ for all ancestors u of v . Thus, decisions in an RDP parametrized by a visibility protocol are always made in light of the updated information regarding detections.

Because K is fixed, the number of visibility protocols is a polynomial in n , but perhaps of high degree. This is because we can enumerate, vertex-by-vertex starting with the root, the possible pairs (z, j) to associate with the nonleaf vertices of T ; the number of possible pairs (z, j) that can be associated with any nonleaf vertex v , given the pair associated with v 's predecessor, is $O(n^2)$. The time to compute the expected length of these RDPs is also a polynomial in n . Thus, the optimization problem in (2) can be solved in polynomial time if we restrict the domain of the optimization problem to visibility protocols T . The subop-

timal protocol solving the restricted problem is denoted \hat{T}^* , and the associated RDP is denoted \hat{p}^* .

The next theorem extends Lemma 2 to the case of an arbitrary number of allowable detections.

THEOREM 3: Let p^* denote an optimal RDP in a mapped hazard field using at most some given $K > 0$ disambiguations, and let \hat{p}^* denote our approximation. Then $El(\hat{p}^*) \leq El(p^*) + 2K\pi\delta$.

The proof of Theorem 3 is analogous to the proof of Lemma 2; each disambiguation creates the possibility of elongation of \hat{p}^* over p^* by $2\pi\delta$ in the exact manner of the single disambiguation of Lemma 2.

Unfortunately, the computation of \hat{T}^* and \hat{p}^* is practical only for small K , and further research into practical, suboptimal protocols is warranted. (In fact, if K is not fixed then the above computation of \hat{T}^* and \hat{p}^* is not polynomial time.)

5. EXAMPLES

5.1. Simulation Experiment

Consider the marked spatial point process \mathcal{G} on $S = [\frac{3}{8}, \frac{5}{8}] \times [-1, 1] \subseteq \mathbb{R}^2$, where the true and false detections are Poisson with parameter value $\lambda = 6$ and $\lambda = 20$, respectively, and the distribution of the probability marks associated with the true and false detections are Beta with parameter values $(6, 2)$ and $(2, 6)$, respectively. We choose $s = (0, 0)$, $d = (1, 0)$ and we suppose $\delta = \frac{1}{8}$, $K = 1$. (See Fig. 3 for a sample realization of this process.)

The simplest question of interest involves the probability that our suboptimal RDP \hat{p}^* is any improvement at all. That is, we want to estimate θ_c , the probability (relative to the underlying process \mathcal{G}) that $El(\hat{p}^*) < \ell^{\mathcal{E}}(q_{s,d,[n]})$ when the disambiguation cost is c . Simulation experiments consisting of 100 Monte Carlo replicates for each of $c = 0$, $c = .025$, and $c = .1$ (which here are 0%, 2.5%, and 10% of $\ell^{\mathcal{E}}(\overline{s}, d) = 1$) yielded the estimates $\hat{\theta}_0 = .70$, $\hat{\theta}_{.025} = .44$, and $\hat{\theta}_{.1} = .18$, indicating a positive expected savings from \hat{p}^* over $q_{s,d,[n]}$ approximately 70%, 44%, and 18% of the time, respectively.

This simulation experiment also affords us the opportunity for a preliminary investigation into the suboptimality of the RDP \hat{p}^* from the optimal RDP p^* . When $c = 0$ the optimal RDP p^* is almost surely shorter than $q_{s,d,[n]}$, unless $q_{s,d,[n]} = s, d$. Among our 100 replicates, we observed 24 where $q_{s,d,[n]} = s, d$, indicating that for 70 of the remaining 76 replicates where optimal RDPs p^* provided an expected savings over $q_{s,d,[n]}$ we also had the suboptimal RDP \hat{p}^* yielding an expected savings over $q_{s,d,[n]}$.

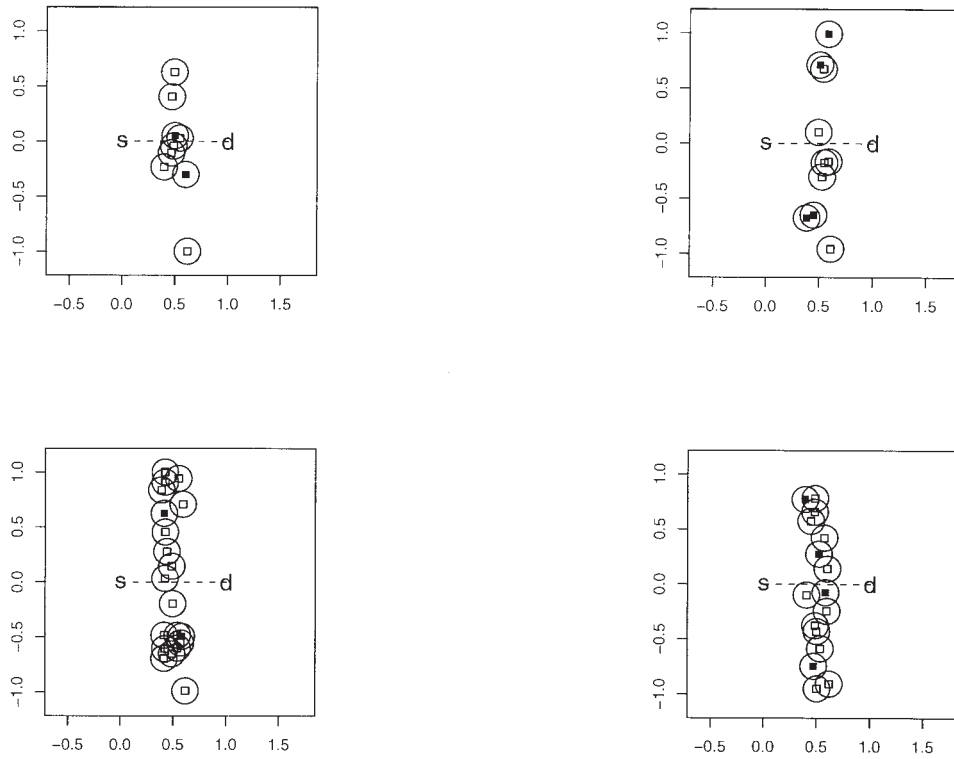


Figure 3. Four representative realizations of the simulation point process. True hazards are depicted as filled squares and false detections are depicted as open squares. Top left: $\ell(q_{s,d,[n]}) = 1.062$ and $\hat{p}^*(0.1) = q_{s,d,[n]}$. Top right: $\ell(q_{s,d,[n]}) = 1.001$ and $\hat{p}^*(0.1) = q_{s,d,[n]}$. Bottom left: $\ell(q_{s,d,[n]}) = 1.196$ and $E[\ell(\hat{p}^*(0.1))] = 1.117$. Bottom right: $\ell(q_{s,d,[n]}) = 2.172$ and $E[\ell(\hat{p}^*(0.1))] = 1.376$.

5.2. Minefield Application

Minefield detection and localization is an important problem currently receiving much attention in the engineering and scientific literature; see, for instance, [16] and references therein. Witherspoon et al. [17] depict the operational concept for minefield reconnaissance via an unmanned aerial vehicle. Multispectral imagery of an area of interest is processed and potential mines are located with a mine detection algorithm. (Holmes et al. [6] present a thorough discussion of the particular mine detection algorithm that we use in the upcoming example.) The detector produces a binary detection map (point pattern) $D(\cdot)$ such that $D(x) = 1$ for all points x at which a mine or minelike object is detected. Categorizing the detections into “true hazards” (mines) and “false detections” (minelike objects, debris, noise, etc.) and considering an operational imperative on the detector to find (nearly) all true hazards, it can be expected that the number of false detections in the map $D(\cdot)$ will be relatively high.

In this subsection we consider the minefield detection risk field presented in Figure 4. Among the 39 detections, 12 are true mines (filled squares) and 27 are false detections (open squares). The marks are posterior probabilities, determined by a post-processing classification rule [9, 12–15], that the

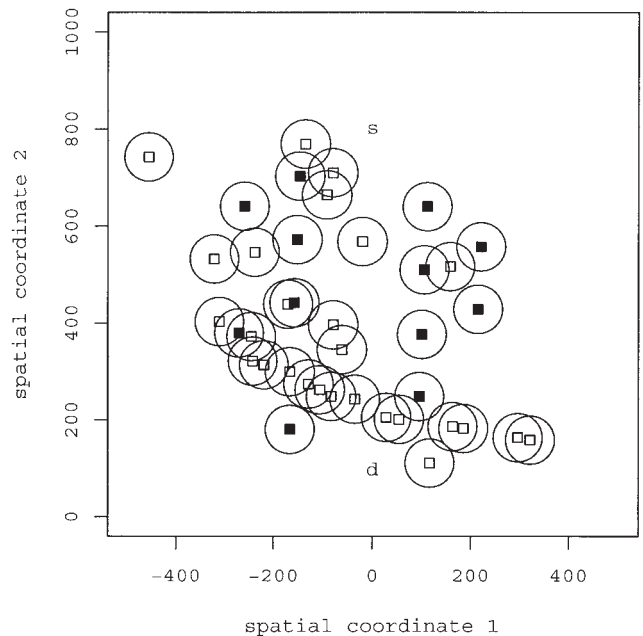


Figure 4. The minefield data set with risk disks. There are 39 detections—12 true mines (filled squares) and 27 false detection (open squares). An example mission requires timely zero-risk transit from source s [at $(0, 800)$ near the top center of the figure] to destination d [at $(0, 100)$ near the bottom center of the figure].

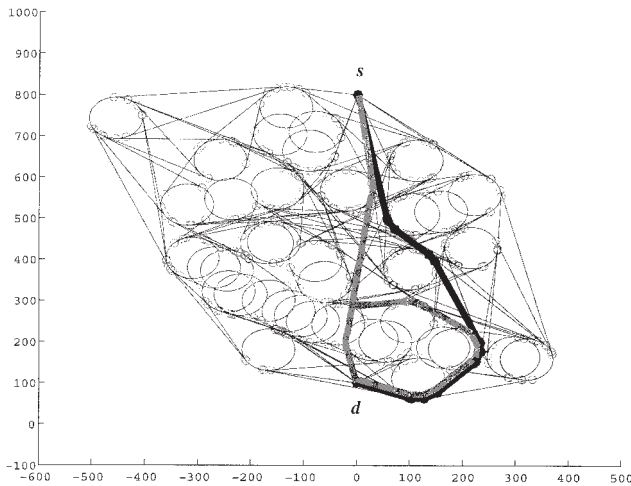


Figure 5. The visibility graph for the minefield data set; the conventional shortest zero-risk path $q_{s,d,[n]}$ is depicted in heavy black, and the random disambiguation path \hat{p}^* for $K = 1$ is depicted in heavy red (gray). The path $q_{s,d,[n]}$ has length $\ell^{\mathcal{E}}(q_{s,d,[n]}) = 978$. The expected length of the RDP \hat{p}^* , which takes the direct route with probability 0.79 and takes the long route with probability 0.21, is less than 978 for disambiguation costs $c \in [0, 185)$.

individual detections obtained from the mine detector are in fact mines.

The path $q_{s,d,[n]}$ shown in heavy black in Figure 5 has length $\ell^{\mathcal{E}}(q_{s,d,[n]}) = 978$. For $K = 1$ disambiguation, the RDP \hat{p}^* , shown in heavy red (gray) in Figure 5 has expected length $E\ell(\hat{p}^*) = 1112\rho + 708(1 - \rho) + c$, where, in this case, $\rho = .21$; thus $E\ell(\hat{p}^*) < \ell^{\mathcal{E}}(q_{s,d,[n]})$ for all $c \in [0, 185)$, and we thus conclude that RDPs are beneficial in this application. We should additionally point out that, although the algorithm does not know which detections are true and which are false, we happen to know here that the detection disambiguated was a false detection, and thus the path actually traversed has length $708 + c$.

It turns out that additional disambiguations are useful in this example, but for $K = 2$ and $K = 3$ the first point of disambiguation is precisely the point of disambiguation in the $K = 1$ case computed above.

6. DISCUSSION

Random disambiguation paths are both theoretically and practically relevant to certain tasks involving path planning under uncertainty; in particular, we have considered the task of minefield traversal. Approximating the optimal RDP can be challenging; a simple procedure has been described here, but additional investigation is demanded. In particular, a dynamic programming approach to the problem may prove useful.

6.1. Stochastic Dynamic Programming

A stochastic dynamic program (SDP) with a properly defined state space can be formulated for our RDP problem. However, it does not currently seem to us that such an SDP can be formulated with a compact representation. The difficulty arises because every state in the SDP state space would need to include the current true/false/ambiguous status of each detection in order for us to be able to determine transition costs with respect to the current potential hazards. In addition, computing such state transition costs would anyway need to be done in a manner similar to that which we have described here. Thus it does not seem that standard SDP techniques can be profitably applied to address our problem more efficiently at this time. The main advantage of our formulation is that we have used visibility graph techniques to explain in a natural way how to compute the underlying costs that would be needed by any SDP formulation. Since our current formulation exploits geometry, this may lead to a more compact state representation and allow more tractable applications of SDP techniques in the future.

6.2. Generalizations

There are many generalizations of our setting which are of both theoretical and practical interest, and we mention but a few: We have restricted our attention to \mathbb{R}^2 , but higher-dimensional spaces are certainly of interest. We use expected length as the goodness criterion, while other goodness criteria may be more situationally appropriate. Multiple hazard types suggest relaxing the assumption that the radii of risk regions are identical, and the required proximity of the sensor to the detection in order to disambiguate may also differ from the radius of the risk region. It is also of interest to generalize to the case of multiple sensors each with its own disambiguation radius and cost, and to the case where disambiguations are not perfect—that is, where each sensor has an associated disambiguation accuracy. The option to neutralize some hazards at additional cost can also be incorporated into the path planning. Also, applications suggest incorporating locational uncertainty into our model. Finally, it is of interest to relax the requirement that the path never enter a risk region, instead allowing for some “tolerable” risk.

ACKNOWLEDGMENTS

The work of CEP was partially supported by Office of Naval Research Grant N00014-04-1-0483, Office of Naval Research Grant N00014-01-1-0011, Office of Naval Research Grant N00024-98-D-8124, and Defense Advanced Research Projects Agency Grant F49620-01-1-0395. The work of CDP was partially supported by Office of Naval

Research Grant N00014-01-01-0964. The authors are grateful to two anonymous referees whose thorough reports and valuable input contributed greatly to the improvement of this manuscript.

REFERENCES

- [1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin, *Network flows: Theory, algorithms, and applications*, Prentice Hall, Englewood Cliffs, NJ, 1993.
- [2] D.P. Bertsekas and J.N. Tsitsiklis, An analysis of stochastic shortest path problems, *Math Oper Res* 16(3) (1991), 580–595.
- [3] A. Briggs, D. Scharstein, and S. Abbott, Reliable mobile robot navigation from unreliable visual cues, *Fourth Int Workshop Alg Found Robot (WAFR 2000)*, Hanover, NH, March 2000.
- [4] T.H. Cormen, C.E. Leiserson, and R.L. Rivest, *Introduction to algorithms*, MIT Press, Cambridge, MA, 1992.
- [5] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational geometry algorithms and applications*, Springer, New York, 1997.
- [6] Q.A. Holmes, C.R. Schwartz, J.H. Seldin, J.A. Wright, and L.J. Wieter, Adaptive multispectral CFAR detection of land mines, *Proc SPIE 2496: Detection Technol Mines Minelike Targets*, Orlando, FL, April 1995, pp. 421–432.
- [7] J.-C. Latombe, *Robot motion planning*, Kluwer Academic, New York, 1991.
- [8] J.S.B. Mitchell, “Shortest paths and networks,” *Handbook of discrete and computational geometry*, J.E. Goodman and J. O’Rourke (Editors), CRC Press, Boca Raton, FL, 1997, pp. 445–466.
- [9] T. Olson, J.S. Pang, and C.E. Priebe, A likelihood—MPEC approach to target classification, *Math Program Ser A* 96 (2002), 1–31.
- [10] C.H. Papadimitriou and M. Yannakakis, Shortest paths without a map, *Theoret Comput Sci* 84 (1991), 127–150.
- [11] S.D. Patek, On partially observed stochastic shortest path problems, *Proc CDC 2001: 40th IEEE Conf Decision Control*, Orlando, FL, December 2001, 5050–5056.
- [12] C.D. Piatko, C.E. Priebe, L.J. Cowen, I.-J. Wang, and P. McNamee, Path planning for mine countermeasures command and control, *Proc SPIE 4394* (2001), 836–843.
- [13] C.D. Piatko, C.P. Diehl, P. McNamee, C. Resch, and I.-J. Wang, Stochastic search and graph techniques for MCM path planning, *Proc SPIE 4742* (2002), 583–592.
- [14] C.E. Priebe, D.Q. Naiman, and L. Cope, Importance sampling for spatial scan analysis: Computing scan statistic p -values for marked point processes, *Comput Statist Data Anal* 35(4) (2001), 475–485.
- [15] C.E. Priebe, J.S. Pang, and T. Olson, Optimizing sensor fusion for classification performance, *Proc CISST ’99 Int Conf*, Las Vegas, NV, June 1999, pp. 397–403.
- [16] D.L. Smith, Detection technologies for mines and minelike targets, *Proc SPIE 2496: Detection Technol Mines Minelike Targets*, Orlando, FL, April 1995, pp. 404–408.
- [17] N.H. Witherspoon, J.H. Holloway, K.S. Davis, R.W. Miller, and A.C. Dubey, The Coastal Battlefield Reconnaissance and Analysis (COBRA) program for minefield detection, *Proc SPIE 2496: Detection Technol Mines Minelike Targets*, Orlando, FL, April 1995, pp. 500–508.