Submitted to *Operations Research*
manuscript 1

Optimal Search with Application to Minimally Invasive Surgery

Farhad Ghassemi
Operations and Logistics Division, Sauder School of Business, University of British Columbia

Steven M. Shechter
Operations and Logistics Division, Sauder School of Business, University of British Columbia

Christopher Y.C. Nguan
Department of Urology, University of British Columbia

Martin L. Puterman
Operations and Logistics Division, Sauder School of Business, University of British Columbia

Motivated by surgical applications, we consider a minimum cost search for a target, whose precise location is uncertain. The region of interest is divided into small grid cells and the searcher may take one of two actions on a cell: 1) cut into the cell, which is quick but bears the risk of damaging the target; or 2) explore the cell, which is time consuming but safely resolves uncertainty regarding whether the target is in that cell. We study two formulations of this problem: the unconstrained search, in which the searcher may visit cells in any order, and the constrained search, in which the searcher may only visit adjacent cells. We prove that the optimal strategy for the unconstrained search problem can be obtained in polynomial time (as a function of the number of cells). We also propose a lookahead strategy for the constrained search problem, and comparisons with the lower bound provided by the unconstrained search problem suggest this strategy performs quite well.

---

1. Introduction

Many surgical procedures can be regarded as optimal searches for a target. For example, surgeons performing minimally invasive surgery (MIS) aim to move their instruments from an initial small incision on the surface of the body (typically near the naval) to a target organ as efficiently as possible. Their goal is to complete the surgery within a reasonable time and to avoid damaging critical structures along the way (including the target organ). While pre-operative images (e.g. CT scans or MRIs) and general anatomical knowledge help the surgeon understand likely locations of these structures beforehand, imperfect imaging techniques and body movement leave uncertainty as to their exact whereabouts (Figure 1). Moreover, unlike traditional surgery with its larger incisions and field of vision, surgeons performing MIS only see what is directly in front of the laparoscope (a tiny camera) inserted into the body.

There are two primary ways during the course of surgery that the surgeon may resolve the

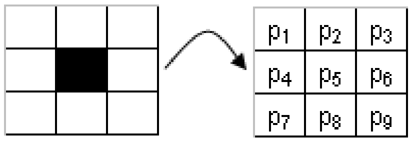

Figure 1 Pre-operative imaging may suggest that the target is in the darkened cell in the middle of the search space, as shown on the left. However, different body position during surgery compared to time of imaging as well as noise around the image results lead to uncertainty of the target's location on the day of surgery. This uncertainty is represented by a probability distribution of possible locations of the target over the search space, as shown on the right.

uncertainty of a target's location: 1) by *cutting* directly into fatty tissue blocking the way; or 2) by slowly peeling away layers of the tissue, which we will refer to as *exploring*, to discover if a critical structure lies just beyond it. A cut is done in near-zero time but carries a large cost if an organ or structure is damaged as a result. On the other hand, exploring takes time but will not cause any damage to critical structures.

Motivated by the above example, we consider the following search problem. We consider a search space discretized into a set of distinct regions and assume that a target occupies one of those regions (from hereon we will refer to the search space as a "grid" and a region as a "cell"). A searcher wishes to minimize the expected cost incurred until locating the target by taking one of two actions after selecting a cell to visit: 1) cutting into the cell, which carries a large cost if the target is in the cell but is otherwise free; or 2) exploring the cell, which carries a moderate cost regardless of what the cell contains. We do not consider a separate cost to move from one cell to another, as this takes a minimal amount of time and is insignificant relative to the time it takes to explore a cell or the consequences of damaging an organ. We assume that the target is certainly detected by both actions if it is present in the cell. The prior information about the location of the target is encapsulated by a probability map (distribution) (Figure 1), which is updated after each action. By weighing the stage costs and analyzing the probability map, the searcher seeks an optimal search strategy which is defined to be a permutation of cells to be visited in order, along with the action to take on each cell upon visiting it. We assume that the actions are always executed as intended. For example, a cut or explore action on a cell will certainly act on that cell and no other cell. In the surgical setting, with the advent of robot-assisted surgery, such precise surgical moves are possible.

We shall consider two formulations of our problem. We refer to the primary formulation as the *constrained search problem*, in which the searcher must start from some initial cell and act on cells adjacent to cells already visited. In other words, the searcher must form paths that emanate from the starting point, which is analogous to a surgeon moving instruments from an initial incision to the target. In the *unconstrained search problem*, the searcher can visit cells in any order. That is, the searcher does not face any physical constraints and can jump around. In order to solve the unconstrained search problem, we demonstrate a number of properties that reduce the number of potential candidates for the optimal strategy from $\mathrm{N}!2^{\mathrm{N}}$ to $\mathrm{N}+1$, where N is the number of cells. Solving the constrained search problem to optimality, however, proves to be challenging. We propose a lookahead strategy for this problem, which uses the expected cost of a greedy strategy to approximate the optimal expected cost. We show the effectiveness of the lookahead strategy by comparing its expected cost to a lower bound for the optimal expected cost, which we obtain from the solution of the unconstrained search problem. The optimal expected cost of the unconstrained search is also used to accelerate the derivation of the lookahead strategy by pruning non-optimal actions through a branch-and-bound routine.

Many search and detection problems try to maximize the probability of detecting a target subject to a constraint on the effort expended by the searcher (Dobbie 1968, Stone 1975). These problems typically arise in searches over large areas (e.g., submarines in the ocean), where there is a chance that the searcher misses the target. Therefore, an overlook probability is often considered, which decreases with the effort spent in searching a cell. Other search problems try to minimize the expected cost of the search until locating the target (Ross 1983, Lossner and Wegener 1982). For example, Ross (Ross 1983) considers "the classical optimal search model"–an unconstrained search problem over N cells that considers a single search action that may be performed on any cell i at a cost of c_i, and which detects the target with probability α_i if it is in that cell (with probability p_i). He shows that it is optimal at each stage to next search cell $i^* = \arg\max \frac{\alpha_i p_i}{c_i}$. Trummel and Weisinger (Trummel and Weisinger 1986) show that a constrained version of that search problem is NP-hard. We also consider an objective of minimizing the expected cost incurred until locating a target, as the surgeon must ultimately find the organ. In the surgical setting, the search cells will be very small and therefore we assume that the target is always detected if it is present in the cell (e.g., $\alpha_i = 1$ for all i). We also consider two search actions and related costs: cut, with an expected cost of $p_i c_H$, and explore, at a fixed cost of c_E. While the consideration of these two actions does not lend itself to an optimal index policy as in Ross's example, we still obtain an efficient algorithm for obtaining the optimal policy of our unconstrained search. However, we suspect our constrained problem is NP-complete and see no efficient way to solve it optimally, leading to our discussion of the lookahead strategy.

Search problems can be formulated as partially observed Markov decision processes (POMDPs) (Smallwood and Sondik 1973, Sondik 1978) in the sense that they are stochastic dynamic programs, where the underlying state is a probability distribution of the target's location over a grid. For a general POMDP, the probability map can evolve to one of potentially many new probability maps as a result of an action (if the target is not located). Therefore, in order to solve a POMDP, a set of decision rules for any possible probability map must be determined. This is the case for many path planning and navigation problems in the robotics literature, in which a robot with noisy sensor measurements seeks to reach a target efficiently (Thrun 2002). For the problem we study here, however, after visiting and acting on a cell i the problem either terminates because the target has been found or the probability map evolves to only one possible new map (which is obtained by taking the remaining cell probabilities and rescaling them to sum to 1, given that $p_i = 0$). In this regard, our problem resembles a deterministic target after visiting at most N cells, the number of possible updates (i.e., states) is still exponential in N (in the case of the unconstrained search, the number of possible updates is 2^{N}, as each subset of cells may remain after some number of steps are taken in a particular search strategy). Therefore, solution by standard dynamic programming methods (e.g., backwards induction) is infeasible. Instead, we take advantage of structural results to solve the problem efficiently (in the case of the unconstrained search) or evaluate the performance of suboptimal strategies (in the case of the constrained search).

Finally, we note that Alterovitz and colleagues (Alterovitz et al. 2005, 2007, 2008) have taken MDP approaches to planning minimally invasive procedures; namely, the optimal steering of a flexible needle through human tissue to reach a target (e.g., to take a biopsy). They consider objectives of minimizing cost (Alterovitz et al. 2005) as well as maximizing probability of success (Alterovitz et al. 2007), formulate an appropriate MDP, and use various solution techniques to obtain an optimal policy for guiding the needle. The key difference between their models and ours is that the motion of the flexible needle is uncertain and they assume the target and obstacle locations are known exactly. Our surgical application deals with rigid instruments, which may be precisely controlled through robot-assisted surgery. Therefore our model assumes no uncertainty in motion, but instead considers uncertainty in the target's location.

The rest of the paper is organized as follows. In Section 2, we formulate the unconstrained search problem and explore structural properties of its solution. In Section 3, we introduce the constrained search problem and discuss the lookahead strategy. In Section 4, we provide the numerical results on the performance of the lookahead strategy. Finally, we provide concluding remarks in Section 5.

2. Formulation and Solution Approach for the Unconstrained Search Problem

In this section, we formulate the unconstrained search problem and prove a number of properties for its optimal strategy. These properties demonstrate that the searcher should first cut through cells that have a low probability of the target being there so that time is not wasted exploring areas where there is very little risk of damaging the target. However, after cutting through some cells and not locating the target, the searcher should switch to the explore action, always executed on the cell that has the highest probability of the target being there.

We consider N cells, labeled from 1 to N, and a target that occupies one of these cells. The *unconstrained search problem* is the solution of

$$\min_{\mathbf{a}\in\mathcal{A}} v_{\mathbf{a}}(\mathbf{x}_1) = \mathbb{E}\left\{\sum_{k=1}^{\mathrm{N}} r(\mathbf{x}_k, a_k)\right\}, \tag{1}$$

with the notation described below and the assumptions made afterwards.

Notation:

1. $\mathbf{a} = \{a_k\}_{k=1}^{\mathrm{N}}$ denotes a sequence of actions to take on some permutation of cells. We call this sequence a search strategy. The action taken at stage k, denoted by a_k, can be either E_i (explore cell i) or C_i (cut through i) for any $i \in \mathcal{A}_k$, where $\mathcal{A}_k$ denotes the set of cells that can be visited at stage k.

2. $\mathcal{A} = \{E,C\}_{\mathcal{A}_1} \times \{E,C\}_{\mathcal{A}_2} \times \cdots \times \{E,C\}_{\mathcal{A}_{\mathrm{N}}}$ denotes the set of feasible strategies ($\{E,C\}_{\mathcal{A}_k}$ indicates the explore or cut actions taken on cells in $\mathcal{A}_k$).

3. $\mathbf{x}_k$ denotes the probability map at stage k, given that the target has not yet been found. Whereas states in MDPs usually have a physical meaning (e.g., inventory level), we choose the state of this problem to be the probability map over the grid of possible target locations. This is in the spirit of POMDP models, in which uncertainty in the underlying state is represented in terms of a probability distribution and the state space is a set of probability distributions. We also note that upon locating the target, the system moves to an absorbing state of 0 cost. We label this state Δ.

4. $r(\mathbf{x}_k, a_k)$ denotes a stage cost incurred in state $\mathbf{x}_k$ if action a_k is taken, with $r(\Delta)$ defined to be 0.

Assumptions: If we assume that $\mathbf{x}_k = \mathbf{p} = (p_1, p_2, \ldots, p_N)$ (where p_i denotes the probability of the target being in cell i),

1. the set of cells that can be visited at stage k (provided the system is not in the absorbing state Δ), is given by

$$\mathcal{A}_k = \{i | p_i \neq 0\}. \tag{2}$$

2. the stage costs are given by

$$r(\mathbf{p}, C_i) = p_i c_H, \tag{3}$$

$$r(\mathbf{p}, E_i) = c_E, \tag{4}$$

$$r(\Delta) = 0 \tag{5}$$

where c_H denotes the cost of accidentally damaging the target as a result of cutting through a cell and c_E ($< c_H$) denotes the cost of exploring. Note that the cost of cutting is an expected value that considers the cost c_H if the target is in cell i (with probability p_i) and the free move cost if it is not there. The explore cost is a fixed cost of c_E regardless of whether the target is in cell i or not.

3. When action E_i or C_i is taken on a cell i, with probability $1 - p_i$ it does not contain the target, and a new probability map results, denoted by $\mathbf{x}_{k+1} = T_i(\mathbf{p})$, where

$$T_i(\mathbf{p})_j = \begin{cases} 0 & \text{if } j = i \text{ and } p_i \neq 1 \\ \frac{p_j}{1-p_i} & \text{if } j \neq i \text{ and } p_i \neq 1 \end{cases} \tag{6}$$

for $j = 1, 2, \cdots, \mathrm{N}$. Note that (6) is the Bayesian update of $\mathbf{p}$. When an action is taken on a cell i, with probability p_i it contains the target, and the system transitions to the absorbing state:

$$T_i(\mathbf{p}) = \Delta. \tag{7}$$

Because the target may be found at any stage, the above problem has a random time horizon. We use a standard approach of formulating such a problem as an infinite horizon MDP (Puterman 1994) as follows:

$$v(\mathbf{p}) = \min_{i \in \mathcal{A}_k} \left\{ \min\{c_E + p_i v(\Delta) + (1-p_i) v(T_i(\mathbf{p})), p_i c_H + p_i v(\Delta) + (1-p_i) v(T_i(\mathbf{p}))\} \right\}, \tag{8}$$

where $v(\mathbf{p})$ denotes the optimal expected cost for a given probability map $\mathbf{p}$. Note that the assumption of perfect detection implies that the random horizon will be of length at most N. Also note that equation (8) explicitly includes the probability p_i of moving to the absorbing state Δ upon finding the target in cell i. However, by the fact that Δ is absorbing and $r(\Delta) = 0$, $v(\Delta) = 0$.

The following proposition introduces a number of properties for the optimal search strategy. Based on these properties, instead of solving the above dynamic program or evaluating all $\mathrm{N}!2^{\mathrm{N}}$ feasible strategies, one can find the optimal strategy by evaluating *at most* $\mathrm{N}+1$ strategies. To prove these properties, we assume that cells 1 to N are labeled such that $p_1 \geq p_2 \geq \cdots \geq p_{\mathrm{N}}$.

Proposition 1. *An optimal strategy for the unconstrained search problem exists and satisfies the following properties:*

1) It consists of a sequence of cut actions followed by a sequence of explore actions.

2) Those cells that are to be explored are done so in ascending order, i.e. cell i must be explored before cell j if $j > i$.

3) Changing the order by which the cells are cut does not change the expected cost of this strategy.

4) The maximum number of cells that are explored is not greater than $\min\{\mathrm{N}, \lceil \frac{c_H}{c_E} \rceil\}$ ($\lceil\ \rceil$ *denotes rounding up to the nearest integer).*

5) The initial probability of the target being in a cell that is cut is always less than the initial probability of the target being in a cell that is explored.

1) We show that the expected cost is never reduced by taking an explore action before a cut action. Consider $E_iC_j\delta$ and $C_jE_i\delta$, where δ is an arbitrary sequence that will be executed after E_iC_j or C_jE_i. We note that

$$v_{E_iC_j\delta}(\mathbf{p}) = c_E + (1-p_i)\left(\frac{p_j}{1-p_i}c_H + (1-\frac{p_j}{1-p_i})v_\delta(T_j(T_i(\mathbf{p})))\right), \tag{9}$$

$$v_{C_jE_i\delta}(\mathbf{p}) = p_jc_H + (1-p_j)\left(c_E + (1-\frac{p_i}{1-p_j})v_\delta(T_i(T_j(\mathbf{p})))\right), \tag{10}$$

from which it follows that

$$v_{E_iC_j\delta}(\mathbf{p}) - v_{C_jE_i\delta}(\mathbf{p}) = p_jc_E \geq 0, \tag{11}$$

considering that $T_i(T_j(\mathbf{p})) = T_j(T_i(\mathbf{p}))$. Accordingly, the cells in an optimal strategy can be partitioned into sets $\mathcal{C}$ and $\mathcal{D}$, where we first cut through the cells in $\mathcal{C}$ and then explore the cells in $\mathcal{D}$. However, there are 2^{N} such partitions and each partition corresponds to $|\mathcal{C}|!|\mathcal{D}|!$ strategies. Therefore, the number of strategies to be evaluated can be still exceedingly large even for small values of N. The other properties that we prove below, however, reduce this number to $\mathrm{N}+1$.

2) Consider $E_iE_j\delta$ and $E_jE_i\delta$, where $j > i$, and hence, $p_i \geq p_j$. We note that

$$v_{E_iE_j\delta}(\mathbf{p}) = c_E + (1-p_i)\left(c_E + (1-\frac{p_j}{1-p_i})v_\delta(T_j(T_i(\mathbf{p})))\right), \tag{12}$$

$$v_{E_jE_i\delta}(\mathbf{p}) = c_E + (1-p_j)\left(c_E + (1-\frac{p_i}{1-p_j})v_\delta(T_i(T_j(\mathbf{p})))\right), \tag{13}$$

from which it follows that

$$v_{E_jE_i\delta}(\mathbf{p}) - v_{E_iE_j\delta}(\mathbf{p}) = (p_i - p_j)c_E \geq 0. \tag{14}$$

3) Consider $C_iC_j\delta$ and $C_jC_i\delta$, where i and j are any arbitrary cells. We note that

$$v_{C_iC_j\delta}(\mathbf{p}) = p_ic_H + (1-p_i)\left(\frac{p_j}{1-p_i}c_H + (1-\frac{p_j}{1-p_i})v_\delta(T_j(T_i(\mathbf{p})))\right), \tag{15}$$

$$v_{C_jC_i\delta}(\mathbf{p}) = p_jc_H + (1-p_j)\left(\frac{p_i}{1-p_j}c_H + (1-\frac{p_i}{1-p_j})v_\delta(T_i(T_j(\mathbf{p})))\right), \tag{16}$$

from which it follows that

$$v_{C_iC_j\delta}(\mathbf{p}) = v_{C_jC_i\delta}(\mathbf{p}) = p_ic_H + p_jc_H + (1-p_i-p_j)v_\delta(T_j(T_i(\mathbf{p}))). \tag{17}$$

4) Let $\mathrm{K} = \lceil \frac{c_H}{c_E} \rceil$, which implies that

$$\mathrm{K}c_E \geq c_H \tag{18}$$

$$(\mathrm{K}-1)c_E < c_H. \tag{19}$$

Consider a scenario where there are K cells left to be visited. To simplify our notation, we assume that these cells are numbered from 1 to K with $p_1 \geq p_2 \geq \cdots \geq\ p_{\mathrm{K}}$ and $\sum_{i=1}^{\mathrm{K}} p_i = 1$. It suffices to show that $E_1E_2\cdots E_{\mathrm{K}}$ cannot be uniquely optimal. We note that

$$\begin{aligned} v_{E_1E_2\cdots E_{\mathrm{K}}}(\mathbf{p}) &= c_E + (1-p_1)c_E + (1-p_1-p_2)c_E + \cdots + (1-p_1-p_2-\cdots-p_{\mathrm{K}-1})c_E \\ &= \left(\mathrm{K} - (\mathrm{K}-1)p_1 - (\mathrm{K}-2)p_2 - \cdots - p_{\mathrm{K}-1}\right)c_E, \end{aligned} \tag{20}$$

$$\begin{aligned} v_{C_{\mathrm{K}}E_1E_2\cdots E_{\mathrm{K}-1}}(\mathbf{p}) &= p_{\mathrm{K}}c_H + (1-p_{\mathrm{K}})c_E + (1-p_{\mathrm{K}}-p_1)c_E + \cdots + (1-p_{\mathrm{K}}-p_1-p_2-\cdots-p_{\mathrm{K}-2})c_E \\ &= p_{\mathrm{K}}c_H + (p_1+p_2+\cdots+p_{\mathrm{K}-1})c_E + (p_2+\cdots+p_{\mathrm{K}-1})c_E + \cdots + p_{\mathrm{K}-1}c_E \\ &= p_{\mathrm{K}}c_H + \left(p_1 + 2p_2 + \cdots + (\mathrm{K}-1)p_{\mathrm{K}-1}\right)c_E, \end{aligned} \tag{21}$$

from which it follows that

$$\begin{aligned} v_{E_1E_2\cdots E_{\mathrm{K}}}(\mathbf{p}) - v_{C_{\mathrm{K}}E_1E_2\cdots E_{\mathrm{K}-1}}(\mathbf{p}) &= \mathrm{K}(1-p_1-p_2-\cdots-p_{\mathrm{K}-1})c_E - p_{\mathrm{K}}c_H \\ &= (\mathrm{K}c_E - c_H)p_{\mathrm{K}} \geq 0, \end{aligned} \tag{22}$$

where the last inequality holds because of (18).

5) Let i be an arbitrarily-chosen cell from $\mathcal{C}$. According to property 3, the order upon which we cut the cells in $\mathcal{C}$ is inconsequential. Hence, we assume that i is the last cell we cut. Also, let j_k be an arbitrarily-chosen cell from $\mathcal{D}$, where j_k is the k-th largest number in $\mathcal{D}$. Consider $C_iE_{j_1}E_{j_2}\cdots E_{j_k}\delta$ and $C_{j_k}E_{j_1}E_{j_2}\cdots E_i\delta$. We note that

$$\begin{aligned} v_{C_iE_{j_1}E_{j_2}\cdots E_{j_k}\delta}(\mathbf{p}) &= p_ic_H + (1-p_i)c_E + (1-p_i-p_{j_1})c_E + \cdots + (1-p_i-p_{j_1}-\cdots-p_{j_{k-1}})c_E \\ &\quad + (1-p_i-p_{j_1}-\cdots-p_{j_{k-1}}-p_{j_k})v_\delta(T_{j_k}(\cdots(T_{j_1}(T_i(\mathbf{p}))))), \end{aligned} \tag{23}$$

$$\begin{aligned} v_{C_{j_k}E_{j_1}E_{j_2}\cdots E_i\delta}(\mathbf{p}) &= p_{j_k}c_H + (1-p_{j_k})c_E + (1-p_{j_k}-p_{j_1})c_E + \cdots + (1-p_{j_k}-p_{j_1}-\cdots-p_{j_{k-1}})c_E \\ &\quad + (1-p_i-p_{j_1}-\cdots-p_{j_{k-1}}-p_{j_k})v_\delta(T_i(\cdots(T_{j_1}(T_{j_k}(\mathbf{p}))))), \end{aligned} \tag{24}$$

from which it follows that

$$v_{C_{j_k}E_{j_1}E_{j_2}\cdots E_i\delta}(\mathbf{p}) - v_{C_iE_{j_1}E_{j_2}\cdots E_{j_k}\delta}(\mathbf{p}) = (p_{j_k} - p_i)(c_H - kc_E) \tag{25}$$

However, from property 4, we know that $c_H \geq kc_E$. Hence, in order for $C_iE_{j_1}E_{j_2}\cdots E_{j_k}$ to be part of an optimal strategy, $p_{j_k} \geq p_i$.

COROLLARY 1. *Based on Proposition 1, there exists an n such that an optimal strategy* $\mathbf{a}^* = \{a_k^*\}_{k=1}^{\mathrm{N}}$ *is given by*

$$a_k^* = \begin{cases} C_{\mathrm{N}-k+1} & \textit{when } k \leq n, \\ E_{k-n} & \textit{when } k > n. \end{cases} \tag{26}$$

As an example, consider $\mathrm{N} = 5$. From property 1, we can conclude that $C_1E_5C_2E_4E_3$ cannot be an optimal strategy but $C_5C_1E_2E_3E_4$ and $C_5C_4E_1E_2E_3$ can be potentially optimal strategies. From property 5, however, we can conclude that it is impossible for $C_5C_1E_2E_3E_4$ to be an optimal strategy.

3. Formulation and Solution Approach for the Constrained Search Problem

In this section, we formulate the constrained search problem and show how we approach this problem. We start by considering the simple case of a constrained search problem on a one-dimensional grid and provide a sufficient condition for the initial probability map that guarantees the optimality of a strategy consisting of a number of cut actions followed by a number of explore actions. We then consider the constrained search problem on a multi-dimensional (e.g., two or three dimensional) grid. We propose a lookahead strategy for this problem where the expected cost of a greedy strategy is used as an approximation for the optimal expected cost. We also use the solution of the unconstrained search problem to accelerate the derivation of this lookahead strategy.

The definition of the constrained search problem is identical to that of the unconstrained search problem. We only need to redefine $\mathcal{A}_k$, the set of cells that can be searched at stage k, as follows:

$$\mathcal{A}_k = \{i | p_i \neq 0 \text{ and there is } j \in \mathcal{N}_i \text{ such that } p_j = 0\}, \tag{27}$$

where $\mathcal{N}_i$ is the set of cells adjacent to cell i. The set $\mathcal{N}_i$ is determined by the geometry of the grid and the region in which the search is carried out. For example, if we consider a two-dimensional grid, $\mathcal{N}_i$ contains the cells above, below, to the right and to the left of cell i (if such cells exist). In other words, $\mathcal{A}_k$ is the set of cells adjacent to cells already visited (see Figure 2), which implicitly makes use of the assumption that moving between cells is free.

Note that we use V throughout this section and remainder of the paper to denote the expected costs related to the constrained search problem (whereas we use v for the unconstrained search problem).

3.1. Constrained Search Problem on One-Dimensional Grids

We first study the constrained search problem on a one-dimensional grid of N cells. Without loss of generality, we assume that the search starts from the leftmost cell and that the cells are labeled

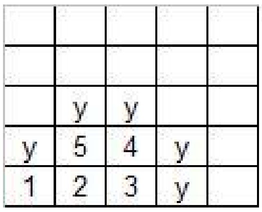

Figure 2 This figure illustrates the description of $\mathcal{A}_k$ in equation (27), which indicates the cells that may be visited next in the search. The searcher visits cells in the order 1-5 as shown (taking either the explore or cut action upon each visit). The "y"s indicate which cells the searcher may visit next, as these are adjacent to cells already visited. For example, while the searcher has just visited cell 5, the zero move cost assumption implies that the searcher may choose to next visit a cell to the right of 3 as well as a cell adjacent to 5.

from 1 to N starting from the leftmost cell and ending with the rightmost cell. Note that for a one-dimensional grid, the order by which the cells are visited is pre-determined, i.e. from left to right or otherwise depending on the initial location of the searcher. Therefore, we only need to determine the type of action taken in each cell. We show that if the initial probability map has the following property, the optimal strategy consists of a sequence of cut actions followed by a sequence of explore actions. This property is adapted from reliability theory (Barlow and Proschan 1965).

Definition 1. Let $h(i) = \frac{p_i}{\sum_{j=i}^{\mathrm{N}} p_j}$ for $i = 1, 2, \cdots, \mathrm{N}$. We say that probability map $\mathbf{p}$ has the *increasing failure rate* (IFR) property if $h(i)$ is non-decreasing in i.

The expression $h(i)$ gives the probability of the target being in cell i given that it is not located to the left of i. This is analogous to the interpretation of the failure rate of an object as being the probability of imminent failure of the object given that it has survived so far. Note that the IFR property is satisfied by both Gaussian and uniform distributions (with the definition modified appropriately for continuous distributions).

Proposition 2. *Suppose the initial probability map on a one-dimensional grid, denoted by* $\mathbf{p}$*, has the IFR property. Then for a search starting from the leftmost cell, there exists a threshold cell* n *such that for* $j < n$*, it is optimal to cut and for* $j \geq n$*, it is optimal to explore.*

C onsider starting from cell 1 and updating $\mathbf{p}$ as the searcher moves to the right. Observe that if we reach cell i and the target is not located, the updated probability in cell i is given by

$$p'_i = \frac{p_i}{\sum_{j=i}^{\mathrm{N}} p_j}. \tag{28}$$

Similarly, if we reach cell $i+1$ and the target is not located, the updated probability in cell $i+1$ is given by

$$p''_{i+1} = \frac{p_{i+1}}{\sum_{j=i+1}^{\mathrm{N}} p_j}. \tag{29}$$

The IFR property implies that $p'_i \le p''_{i+1}$. Therefore, if one comes upon a cell i for which $c_E \le p'_i c_H$ (that is, an explore action is optimal in cell i), then this action is optimal for cell $i+1$ and all cells to the right of i.

We next consider the effect of reducing the uncertainty around the target's location. Presumably, there is a benefit to improved imaging/information that produces smaller variance. We explore this in the context of uniformly distributed probability maps with different numbers of cells. Note that the variance of a discrete uniform distribution over N cells is given by $\frac{\mathrm{N}^2-1}{12}$, so a smaller region will have smaller location variance.

The following proposition shows that starting the search in the cell immediately to the left of a larger uncertainty region has a greater expected cost than starting adjacent to a smaller uncertainty region. In the proposition, we use $V(k,\mathbf{p})$ to indicate the optimal expected cost for the *initial* probability map $\mathbf{p}$ when the searcher has reached cell k without locating the target. Note that knowing the current location of the searcher and the initial probability map suffice to specify the current probability map in the one-dimensional constrained search problem.

Proposition 3. *Consider two initial probability maps, denoted by* $\mathbf{p}$ *and* $\mathbf{q}$*, that are defined on a one-dimensional grid of* N *cells as follows:*

$$p_i = \begin{cases} \frac{1}{\mathrm{N}_A} & \text{if } i \le \mathrm{N}_A, \\ 0 & \text{if } i > \mathrm{N}_A, \end{cases} \qquad q_i = \begin{cases} \frac{1}{\mathrm{N}_B} & \text{if } i \le \mathrm{N}_B, \\ 0 & \text{if } i > \mathrm{N}_B, \end{cases} \tag{30}$$

where $\mathrm{N}_A > \mathrm{N}_B$ *and* $1 \le \mathrm{N}_A, \mathrm{N}_B \le \mathrm{N}$. *Then,* $V(0,\mathbf{p}) \ge V(0,\mathbf{q})$.

F irst note that by the Bayesian updates, $V(0,\mathbf{q}) = V(\mathrm{N}_A - \mathrm{N}_B, \mathbf{p})$. Then, it suffices to show that $V(k,\mathbf{p})$ is decreasing in k for $0 \le k \le \mathrm{N}$. To this end, we note that $V(k,\mathbf{p})$ satisfies

$$V(k,\mathbf{p}) = \min\left\{ c_E + \frac{\mathrm{N}_A - k - 1}{\mathrm{N}_A - k} V(k+1,\mathbf{p}), \frac{c_H}{\mathrm{N}_A - k} + \frac{\mathrm{N}_A - k - 1}{\mathrm{N}_A - k} V(k+1,\mathbf{p}) \right\} \tag{31}$$

for $0 \le k \le \mathrm{N} - 1$. Therefore, we face the following two possibilities:

1) When the explore action is optimal, we need to check if

$$c_E + \frac{\mathrm{N}_A - k - 1}{\mathrm{N}_A - k} V(k+1,\mathbf{p}) \ge V(k+1,\mathbf{p}), \tag{32}$$

which holds if

$$(\mathrm{N}_A - k) c_E \ge V(k+1,\mathbf{p}). \tag{33}$$

Now consider a strategy of exploring all cells. Following the argument used to derive (20), we can see that the expected cost of this strategy from cell $k+1$ is less than the left-hand side of (33).

Hence, the expected cost of the optimal strategy from cell $k+1, V(k+1,\mathbf{p})$, is also less than the left-hand side of (33).

2) When the cut action is optimal, we need to check if

$$\frac{c_H}{\mathrm{N}_A - k} + \frac{\mathrm{N}_A - k - 1}{\mathrm{N}_A - k} V(k+1,\mathbf{p}) \geq V(k+1,\mathbf{p}), \tag{34}$$

which holds if

$$c_H \geq V(k+1,\mathbf{p}) \tag{35}$$

Consider a strategy of cutting all cells. One can expand (17) or use the intuitive fact that cutting through all cells will certainly end up damaging the target to reason that the expected cost of this strategy starting from any cell equals c_H. Hence, the expected cost of the optimal strategy from cell $k+1, V(k+1,\mathbf{p})$, is also less than c_H.

3.2. Constrained Search Problem on Multi-Dimensional Grids

We next consider the constrained search problem on a multi-dimensional grid. Similar to the unconstrained search problem, here we must determine both the order by which the cells are visited and the type of action taken on each cell. The properties that we obtained to guarantee a quick exact optimal solution for the unconstrained search problem, however, do not hold in the unconstrained context. Therefore, we propose a lookahead strategy that is shown to be very close to optimality based on the results presented in Section 4. The general idea of a s-stage lookahead policy is to approximate the optimal expected cost that may arise after s stages and then use that value in choosing optimal actions for s stages prior to that (Bertsekas 2000). As an example, consider a one-stage lookahead policy, where only the action for the next immediate stage is determined and the optimal expected cost of the remaining stages is approximated. Given that $\mathbf{x}_k = \mathbf{p}$, the k-th action in this strategy, denoted by a_k^1, is the one that solves

$$\min_{i \in \mathcal{A}_k} \left\{ \min\{c_E + (1-p_i)\hat{V}(T_i(\mathbf{p})), p_i c_H + (1-p_i)\hat{V}(T_i(\mathbf{p}))\} \right\}, \tag{36}$$

where $\mathcal{A}_k$ is given by (27) and $\hat{V}(T_i(\mathbf{p}))$ is an approximation for $V(T_i(\mathbf{p}))$. Similarly, the two-stage lookahead strategy, denoted by a_k^2, can be obtained by solving

$$\begin{aligned} \min\nolimits_{i \in \mathcal{A}_k} \Bigg\{ \min \Bigg\{ & c_E + (1-p_i)(\min\nolimits_{j \in \mathcal{A}_{k+1}} \left\{ \min\{c_E + (1-\tfrac{p_j}{1-p_i})\hat{V}(T_j(T_i(\mathbf{p}))), \tfrac{p_j}{1-p_i} c_H + (1-\tfrac{p_j}{1-p_i})\hat{V}(T_j(T_i(\mathbf{p})))\} \right\}), \\ & p_i c_H + (1-p_i)(\min\nolimits_{j \in \mathcal{A}_{k+1}} \left\{ \min\{c_E + (1-\tfrac{p_j}{1-p_i})\hat{V}(T_j(T_i(\mathbf{p}))), \tfrac{p_j}{1-p_i} c_H + (1-\tfrac{p_j}{1-p_i})\hat{V}(T_j(T_i(\mathbf{p})))\} \right\}) \Bigg\} \Bigg\}. \end{aligned} \tag{37}$$

We denote the s-stage lookahead strategy by $\mathbf{a}^s = \{a_k^s\}_{k=1}^{\mathrm{N}}$ and its expected cost for the initial probability map $\mathbf{p}$ by $V_{\mathbf{a}^s}(\mathbf{p})$. The quality of a lookahead strategy depends on how the optimal expected cost is approximated. Ideally the approximation is both a good one and easily computable. We consider the expected cost of a greedy strategy for this purpose. Given $\mathbf{x}_k = \mathbf{p}$, the k-th action in this strategy can be expressed as

$$g_k = L_{i^*}, \tag{38}$$

where

$$i^* = \arg\max_{i \in \mathcal{A}_k} p_i, \tag{39}$$

$$L = \arg\min\{c_E, p_{i^*} c_H\}. \tag{40}$$

In other words with the greedy strategy, we always visit the cell that has the highest probability of the target being there and choose the type of action that incurs the lowest stage cost. We denote the greedy strategy by $\mathbf{g} = \{g_k\}_{k=1}^{\mathrm{N}}$ and its expected cost for the initial probability map $\mathbf{p}$ by $V_{\mathbf{g}}(\mathbf{p})$.

Naturally one would expect that as the number of stages increases, the effect of the approximation becomes less significant and the lookahead strategy converges to the optimal strategy. However, it also becomes computationally more difficult to find the lookahead strategy as the number of stages increases. In order to accelerate the derivation of the lookahead strategy, we can take advantage of our results for the unconstrained search problem and prune some of the non-optimal actions for the lookahead strategy. By noting that the expected cost of the unconstrained problem provides a lower bound for the expected of the constrained problem, we can achieve this goal. Consider the case that we must determine the k-th action of a s-stage lookahead strategy, where $\mathbf{x}_k = \mathbf{p}$ and $s > 1$. For a particular cell i where $i \in \mathcal{A}_k$, let δ denote any arbitrary strategy that can be taken from stage $k+1$ conditioned on that the k-th action is taken on cell i. The following then always hold:

$$V_{a_k\delta}(\mathbf{p}) \geq \min\{c_E, p_i c_H\} + (1 - p_i) v(T_i(\mathbf{p})). \tag{41}$$

In other words, the right-hand side of (41) provides a lower bound for $V_{a_k\delta}(\mathbf{p})$. Let $\underline{V}_{a_k\delta}(\mathbf{p})$ denote this lower bound. Now, consider cell j where $j \in \mathcal{A}_k$. let β denote any arbitrary strategy that can be taken from stage $k+1$ conditioned on that the k-th action is taken on cell j. The following then always hold:

$$\min\{c_E, p_j c_H\} + (1 - p_j) V_{\mathbf{g}}(T_j(\mathbf{p})) \geq V_{a_k\beta}(\mathbf{p}). \tag{42}$$

Hence, the left-hand side of (42) provides an upper bound for $V_{a_k\beta}(\mathbf{p})$. Let $\overline{V}_{a_k\beta}(\mathbf{p})$ denote this upper bound. If we have

$$\underline{V}_{a_k\delta}(\mathbf{p}) \geq \overline{V}_{a_k\beta}(\mathbf{p}), \tag{43}$$

then one does not need to consider cell i and all subsequent strategies that are initiated from this cell in order to find the optimal action at stage k.

4. Numerical Results

In this section, we study the performance of the lookahead strategy through numerical examples. In each example, we consider a two-dimensional grid of 10×10 (i.e., $\mathrm{N} = 100$), which means the number of feasible strategies for the unconstrained search problem is in the order of 10^{188}. For the constrained search problem, we always assume that the initial location of the searcher is the cell at the lower-left corner of the grid. In the first numerical example, we assume that the initial probability map is extracted from a Gaussian distribution, which is often used to represent uncertainty around a measurement (in this case, the pre-operative images). We compute the initial probability of the target being in each cell by integrating the volume of the Gaussian distribution over the area of the cell and scaling these probabilities to add up to one. In the second numerical example, we assume that the initial probability map is extracted from a mixture of Gaussian distributions. Although it is unlikely that one encounters such a probability distribution in surgical applications, we study this probability distribution to demonstrate the effectiveness of the lookahead strategy regardless of the choice of the initial probability map.

4.1. Performance Study When the Initial Probability Map is Extracted from a Gaussian Distribution

We denote a Gaussian distribution by $N(\boldsymbol{\mu}, \mathbf{C})$, where $\boldsymbol{\mu}$ is the mean and $\mathbf{C}$ is the covariance. We assume that the mean is in the center of the grid, i.e. $\boldsymbol{\mu} = [5, 5]$. We also consider $c_H = 50$ and $c_E = 1$. According to Proposition 1, these values ensure that the maximum number of cells that are explored in the optimal strategy for the unconstrained search problem is 50. Figure 3 illustrates the optimal strategy for the unconstrained search problem, the greedy strategy for the constrained search problem and the 1-stage lookahead strategy for the constrained search problem, when the covariance of the Gaussian distribution is $9\mathbf{I}$ ($\mathbf{I}$ denotes the identity matrix). In other words, the x and y-coordinates of the target's location are independent of each other and the standard deviation of the location in each direction is 3. From Figure 3(a), we can see that the optimal strategy for the unconstrained search problem consists of first cutting through an outer ring of lowest probability cells for the first 49 moves before jumping to the highest probability cells in the center at move 50, and then exploring outward from those cells. From Figure 3(b), we can see that in the greedy strategy for the constrained search problem, the searcher first reaches the center via the shortest path by cutting through the cells and then explores the remaining cells. Figure 3(c) suggests that the lookahead strategy for the constrained search problem is very similar to the optimal strategy for the unconstrained search problem in terms of the cells that are explored or cut. The two strategies, however, differ in the order by which the cells are visited.

Note that the assumption that it is free to move between cells already visited leads to some jumps between non-adjacent cells in the constrained version of the problem (e.g., see the explore actions taken at moves 59 and 60 in Figure 3(c)). This kind of behavior could still arise even if we included a positive cost of moving between cells, c_M, where $c_M << c_E < c_H$. While surgeons may quickly move through areas already explored, there are limits to how often and when they will do this. We discuss how we might adapt our model to consider these factors further in Section 5.

Figure 4 compares the performance of the three strategies as a function of the initial uncertainty about the location of the target. We consider $\sigma^2\mathbf{I}$ as the covariance of the Gaussian distribution and quantify uncertainty by σ (the mean and cost parameters remain as before). Figure 4(a) represents the expected costs of the optimal strategy for the unconstrained search problem and the expected cost of the greedy strategy and the 1-stage lookahead strategy for the constrained search problem. From the figure, we can see that as σ increases, the expected cost of all strategies increases. Figure 4(b) represents the relative difference in the expected cost of the two latter strategies with respect to the expected cost of the optimal strategy for the unconstrained search problem. More accurately, this figure represents $\frac{V_{a^1}(\mathbf{p})-v(\mathbf{p})}{v(\mathbf{p})}$ and $\frac{V_g(\mathbf{p})-v(\mathbf{p})}{v(\mathbf{p})}$. From the figure, it appears that the lookahead strategy always outperforms the greedy strategy. It also appears that the expected cost of the lookahead strategy converges to the expected cost of the optimal strategy for the unconstrained search problem as σ increases. In general, we expect that the optimal expected cost of the constrained search problem converges to the optimal expected cost of the unconstrained search problem as σ increases. This can be explained by the fact that when σ is small, the probability mass is highly concentrated in the center and the expected cost of any strategy is mainly determined by stage costs incurred in the central cells. Therefore, if two strategies are different in the timing of their visits to these cells, their expected costs will be different. On other hand, when σ is large, the probability mass is more evenly distributed. Therefore, at any given stage in the constrained search problem, it is likely to find a cell that can be visited and has a probability close to the probability of the cell that is visited by the optimal strategy for the unconstrained search problem. Therefore, the expected costs of two strategies become closer.

Figure 4(b) also illustrates that the expected cost of the lookahead strategy is very close to the expected cost of the optimal strategy for the unconstrained search problem at least when σ is large.

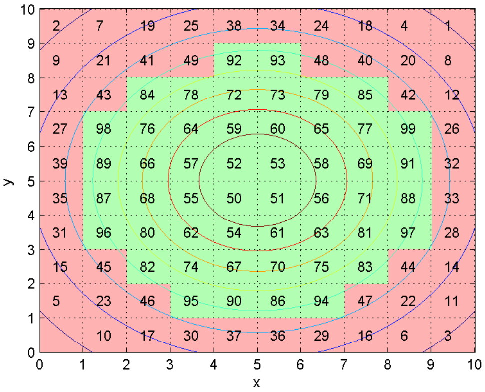

(a)Optimal strategy for the unconstrained search problem

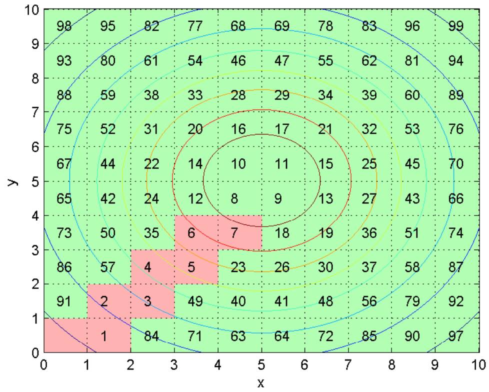

(b)Greedy strategy for the constrained search problem

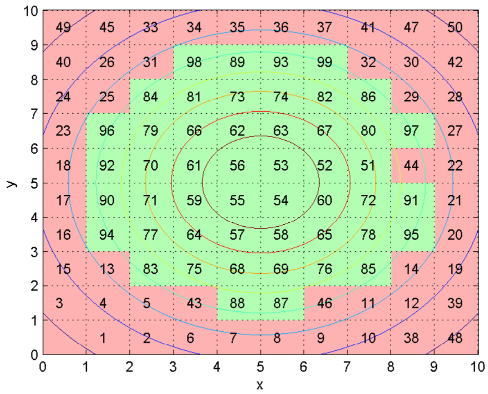

(c)lookahead strategy for the constrained search problem

Figure 3 The optimal strategy for the unconstrained search problem, the greedy strategy for the constrained search problem and the lookahead strategy for the constrained search problem when the initial probability map is extracted from a Gaussian distribution with the mean of $[5,5]$ and the covariance of $9\mathbf{I}$. The red (dark-colored) cells are cut and the green (light-colored) cells are explored. The number in each cell denotes the order by which the cell is visited. The contours displays the isolines of the distribution.

Therefore, it can be argued that the lookahead strategy is very close to the optimal strategy for the constrained search problem. As another indicator that this strategy is, in fact, close to optimality when the initial probability map is extracted from a Gaussian distribution, we also observed that the expected cost of the lookahead strategy does not change with the number of stages in the lookahead strategy over the whole range assigned to σ. In Section 4.2, however, we will see that the expected cost of the lookahead strategy changes with number of stages when the initial probability map is extracted from a mixture of Gaussian distributions.

Figure 5 compares the performance of the three strategies as a function of the hit to explore cost

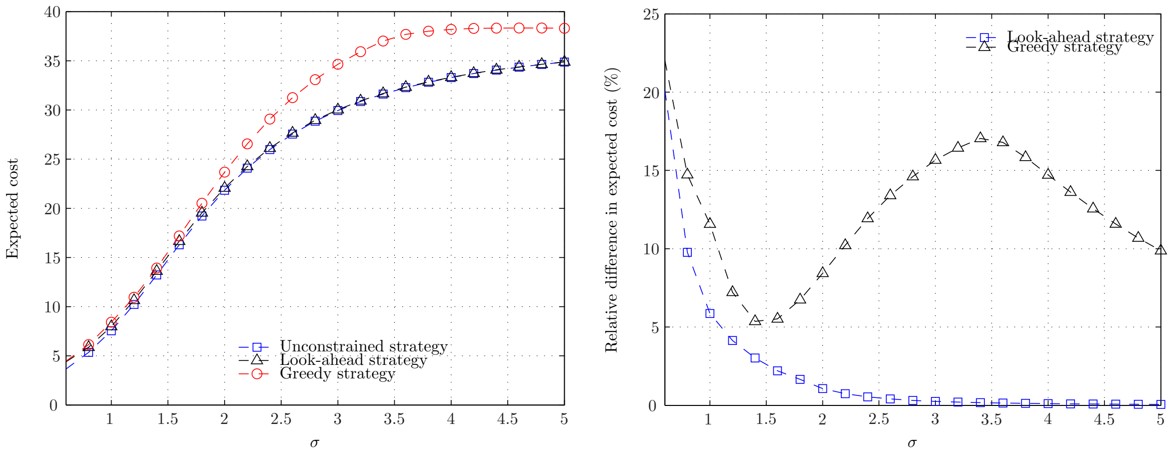

(a)Expected costs of different strategies (b)Relative difference in expected costs

Figure 4 The performance of different strategies as a function of the uncertainty in the initial probability map when $\frac{c_H}{c_E} = 50$. Figure 4(a) represents the absolute expected cost of these strategies. Figure 4(b) represents the relative difference in the expected costs with respect to the expected cost of the optimal strategy for the unconstrained search problem.

ratio ($\frac{c_H}{c_E}$) when $\boldsymbol{\mu} = [5,5]$ and $\mathbf{C} = 9\mathbf{I}$. Figure 5(a) suggests that the expected costs of all strategies increase as $\frac{c_H}{c_E}$ increases. Figure 5(b) also confirms that the expected cost of the lookahead strategy remains very close to the expected cost of the optimal strategy for the unconstrained search problem regardless of the values chosen for c_H and c_E. We can see in the figure that the difference between the expected cost of the lookahead strategy and the optimal strategy for the unconstrained search problem increases as $\frac{c_H}{c_E}$ increases. This behavior can be justified by noting that as $\frac{c_H}{c_E}$ increases, the number of cells that are explored increases and the jump that happens in the optimal strategy for the unconstrained search problem from the outer cell to the central cells becomes wider. Hence, the difference between the two strategies for the constrained and unconstrained search problems widens.

4.2. Performance Study When the Initial Probability Map is Extracted from a Mixture of Gaussian Distributions

For this numerical example, we consider an initial probability map that is given by $\frac{1}{2}N(\boldsymbol{\mu}_1, \mathbf{C}_1) + \frac{1}{2}N(\boldsymbol{\mu}_2, \mathbf{C}_2)$. Figure 6 illustrates the optimal strategy for the unconstrained search problem, the greedy strategy for the constrained search problem and the 2-stage lookahead strategy for the constrained search problem when $\boldsymbol{\mu}_1 = [7,7]$, $\boldsymbol{\mu}_2 = [2,2]$ and $\mathbf{C}_1 = \mathbf{C}_2 = 4\mathbf{I}$. From the figure, we can see that, similar to the previous numerical example, there is a large overlap between the areas that are explored or cut in both the optimal strategy for the unconstrained search problem and the lookahead strategy.

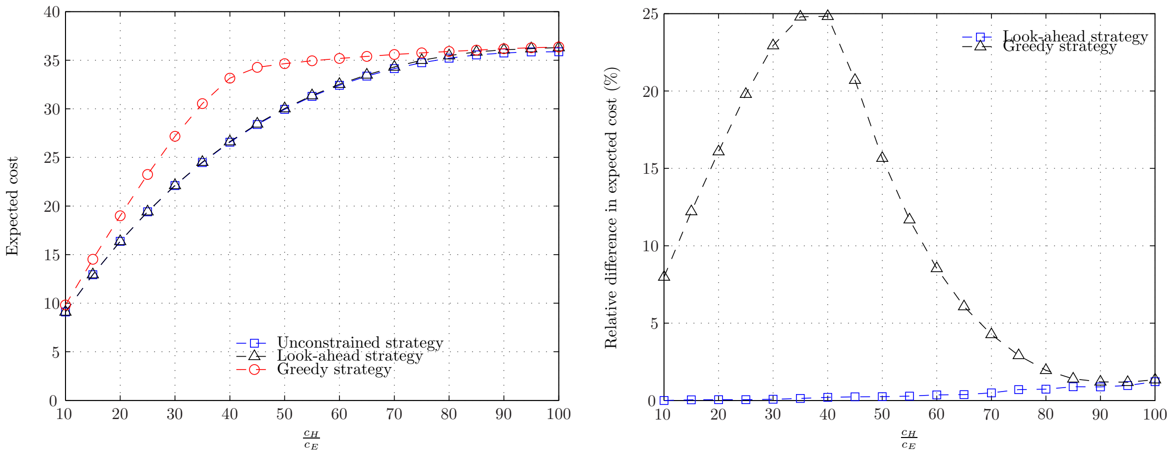

(a)Expected costs of different strategies (b)Relative difference in expected costs

Figure 5 The performance of different strategies as a function of $\frac{c_H}{c_E}$ when $\boldsymbol{\mu} = [5,5]$ and $\mathbf{C} = 9\mathbf{I}$. Figure 5(a) represents the absolute expected cost of these strategies. Figure 5(b) represents the relative difference in the expected costs with respect to the expected cost of the optimal strategy for the unconstrained search problem.

Figure 7 compares the performance of the three strategies as a function of the bi-modality in the initial probability map when $\frac{c_H}{c_E} = 50$. We assume that the initial probability map is given by $\frac{1}{2}N([5,5],1\mathbf{I}) + \frac{1}{2}N(m[1,1],1\mathbf{I})$ and change the value of m from 1.5 to 8.5. Hence, m can be considered as a measure of bi-modality in the distribution (When $m = 5$, we have a uni-modal distribution). Figure 7(a) represents the expected cost of the optimal strategy for the unconstrained search problem and the expected cost of the greedy strategy and the 2-stage lookahead strategy for the constrained search problem. From the figure, we can see that, in general, the expected cost of all strategies decreases as bi-modality decreases. Figure 7(b) represents the relative difference in the expected cost of the two latter strategies with respect to the expected cost of the optimal strategy for the unconstrained search problem. As before, these relative differences are given by $\frac{V_{a^2}(\mathbf{p})-v(\mathbf{p})}{v(\mathbf{p})}$ and $\frac{V_g(\mathbf{p})-v(\mathbf{p})}{v(\mathbf{p})}$. From the figure, it appears that, similar to the previous numerical example, the lookahead strategy always outperforms the greedy strategy.

In Figure 8, we investigate how the expected cost of the lookahead strategy changes with the number of stages used to obtain this strategy. We investigate this relation where various other parameters of the problem are changed one by one. For the base case, we consider the following values: $\boldsymbol{\mu}_1 = [8,8]$, $\boldsymbol{\mu}_2 = [2,2]$, $\mathbf{C}_1 = \mathbf{C}_2 = 4\mathbf{I}$ and $\frac{c_H}{c_E} = 50$. In the second case, we change $\frac{c_H}{c_E}$ to 100 while maintaining the values of the other parameters. In the third case, we change $\boldsymbol{\mu}_2$ to $[6,6]$. Finally in the fourth case, we change $\mathbf{C}_1$ and $\mathbf{C}_2$ to $16\mathbf{I}$. Figure 8 shows that the expected cost, in general, decreases as the number of stages increases. In the third and fourth cases, these

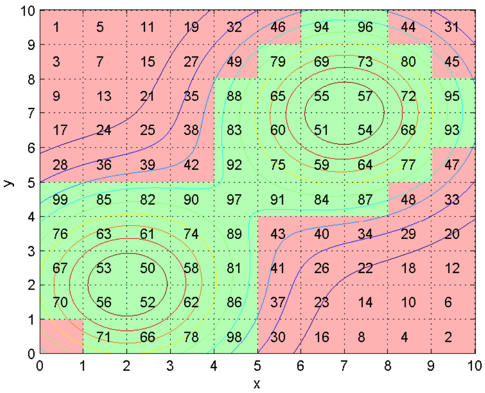

(a)Optimal strategy for the unconstrained search problem

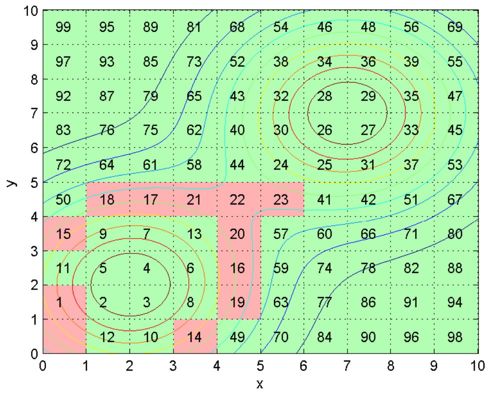

(b)Greedy strategy for the constrained search problem

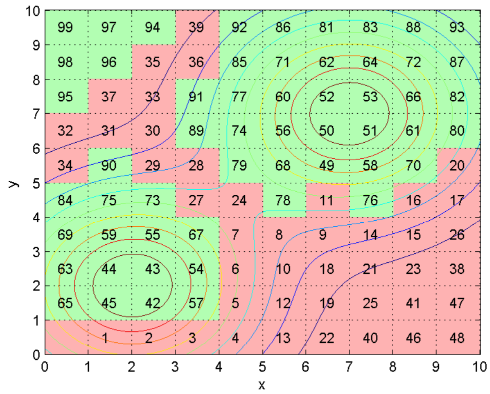

(c)lookahead strategy for the constrained search problem

Figure 6 The optimal strategy for the unconstrained search problem, the greedy strategy for the constrained search problem and the lookahead strategy for the constrained search problem when the initial probability map is extracted from a mixture of Gaussian distributions given by $\frac{1}{2}N([7,7],4\mathbf{I})+\frac{1}{2}N([2,2],4\mathbf{I})$. The red (dark-colored) cells are cut and the green (light-colored) cells are explored. The number in each cell denotes the order by which the cell is visited. The contours displays the isolines of the distribution.

changes are, in fact, very small. In the figure, the relative differences are generated with respect to the expected cost of the 4-stage lookahead strategy in each case, i.e. the vertical axis represents $\frac{V_{a^s}(\mathbf{p})-V_{a^4}(\mathbf{p})}{V_{a^4}(\mathbf{p})}$, where $s=1,2,3,4$.

5. Discussion

Motivated by minimally invasive surgery procedures, we have proposed a search problem that captures two key aspects of these operations: 1) uncertainty in an organ's exact location, and 2) the types of decisions surgeons face to help resolve this uncertainty as they search for the

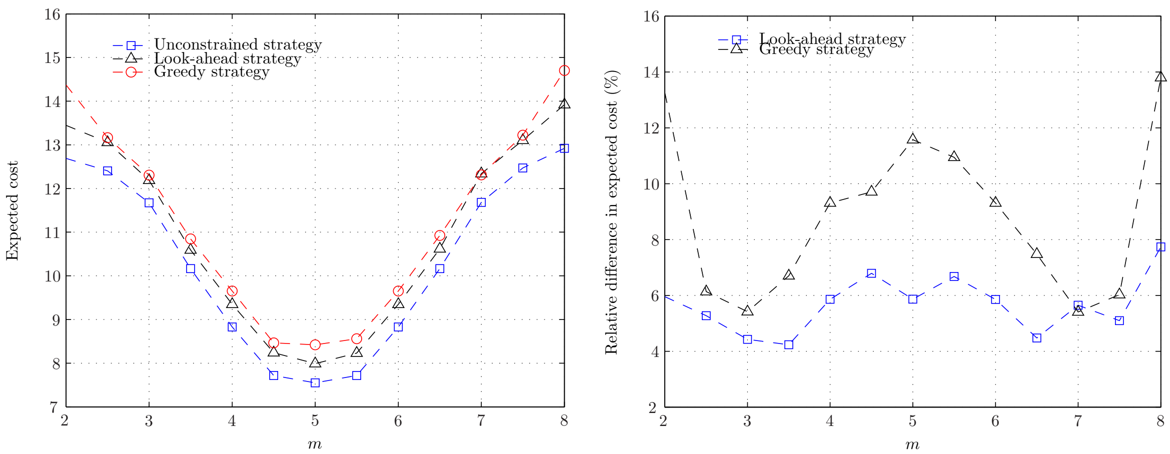

(a)Expected costs of different strategies (b)Relative difference in expected costs

Figure 7 The performance of different strategies as a function of m when the initial probability map is given by $\frac{1}{2}N([5,5],1\mathbf{I})+\frac{1}{2}N(m[1,1],1\mathbf{I})$ and $\frac{c_H}{c_E}=50$. Figure 7(a) represents the absolute expected cost of these strategies. Figure 7(b) represents the relative difference in the expected costs with respect to the expected cost of the optimal strategy for the unconstrained search problem.

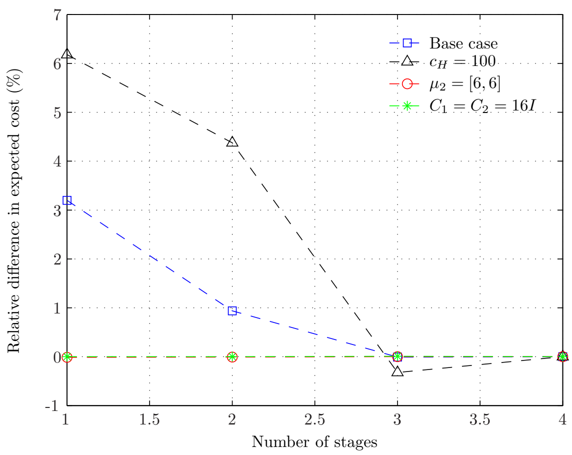

Figure 8 The performance of lookahead strategies as a function of the number of stages used to obtain this strategies. As a base case, we consider $\boldsymbol{\mu}_1=[8,8]$, $\boldsymbol{\mu}_2=[2,2]$, $\mathbf{C}_1=\mathbf{C}_2=4\mathbf{I}$ and $\frac{c_H}{c_E}=50$. We then change the value of one parameter in each case and observe how the expected cost changes as the function of the number of stages with respect to the expected cost of the 4-stage lookahead strategy.

organ: a quick cut vs. a slower explore action. We first formulated an unconstrained version of the problem which we can solve efficiently and whose solution provides a lower bound on the constrained search (which is a more realistic search model in the surgical setting). The lower bound is useful both for comparison to the expected costs of the proposed lookahead strategies as well as for implementing a branch-and-bound routine within our derivation of those strategies. Our numerical results demonstrated that the lookahead strategy performs well under several scenarios.

In addition to the experiments presented in Section 4, we also tested the lookahead strategy on randomly generated probability maps and always found the expected cost to be within 10% of the lower bound provided by the solution of the unconstrained search problem.

It is worthwhile to mention another strategy for the constrained search problem, which may be particularly useful when the initial probability map is extracted from a Gaussian distribution. The idea behind this strategy is to try to mimic the optimal strategy for the unconstrained search problem as closely as possible. Referring back to Figure 3(a), we see that the optimal unconstrained solution cuts around an outer ring before jumping to the highest probability cells in the center and then exploring from the inside region out. By the symmetry of the distribution, many of these explore moves are performed sequentially on cells adjacent to already explored cells. Furthermore, as noted by Proposition 1, part 3, the outer region that is cut can be done so in any order and therefore can be rearranged to cut through cells adjacent to already cut cells. Therefore, the constrained problem, which must act on cells adjacent to already visited cells, can for the most part mimic the unconstrained solution in the cut and explore regions, respectively. The only place where it will differ is when the unconstrained solution jumps from the cut region to the center (mode) of the Gaussian probability map, which the constrained solution cannot do. Instead, a shortest path route may be taken from the last cut cell to the center cell. In every experiment we ran, the lookahead strategy outperformed the mimic strategy; however, their expected costs were relatively close and the computational time needed for the mimic strategy was shorter. On the other hand, it will become more difficult to mimic the unconstrained solution as the problem complexity increases, whereas the lookahead strategy is general enough to be a reasonable approach to a wider class of problems.

During the course of an operation, surgeons might approach an organ from different directions and locate different regions of the organ. If it is deemed that locating a particular region is difficult, they may attempt to approach the organ from another angle and locate a different area along its surface. Our current formulation only considers a target represented by a single cell. As a step to increase the applicability our work to surgery, we plan to investigate locating a target represented by multiple cells so that different regions of the target might be reached as part of the objective. The possibility of changing approach directions also relates to our earlier discussion in Section 4 regarding the jumps between non-adjacent cells exhibited by the constrained search solutions. To limit the frequency with which this may occur, we could impose a constraint such as a minimum number of visits that must be made on adjacent cells before moving back through cells already visited. Finally, a key input to our model is the probability map describing the possible locations of the target organ, and we are currently investigating ways to derive clinically-based location distributions.

Acknowledgments
The authors would like to thank Midhun Salim, supported by the MITACS Globalink Internship Program, for his help with Matlab programming and the numerical experiments. Steven Shechter was supported by NSERC Discovery Grant 341415-07. Martin Puterman and Farhad Ghassemi were supported by NSERC Discovery Grant 5527-07.

References

Alterovitz, R., M. Branicky, K. Goldberg. 2008. Motion planning under uncertainty for image-guided medical needle steering. *International Journal of Robotics Research* **7**(11-12) 1361–1374.

Alterovitz, R., A. Lim, K. Goldberg, G. S. Chirikjian, A. M. Okamura. 2005. Steering flexible needles under Markov motion uncertainty. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* 120–125.

Alterovitz, R., T. Simeon, K. Goldberg. 2007. The stochastic motion roadmap: A sampling framework for planning with Markov motion uncertainty. *Robotics: Science and Systems III, June 27-30* 1–8.

Barlow, R. E., F. Proschan. 1965. *Mathematical Theory of Reliability*. John Wiley & Sons.

Bertsekas, D. P. 2000. *Dynamic Programming and Optimal Control*. Athena Scientific.

Dobbie, J. M. 1968. A survey of search theory. *Operations Research* **16**(3) 527–537.

Lossner, U., I. Wegener. 1982. Discrete sequential search with positive switch cost. *Mathematics of Operations Research* **7**(3) 426–440.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience.

Ross, S. M. 1983. *Introduction to Stochastic Dynamic Programming*. Academic Press.

Smallwood, R.D., E.J. Sondik. 1973. The optimal control of partially observable Markov decision processes over a finite horizon. *Operations Research* **21** 1071–1088.

Sondik, E.J. 1978. The optimal control of partially observable Markov processes over the infinite horizon: discounted cost. *Operations Research* **26** 282–304.

Stone, L. D. 1975. *Theory of Optimal Search*. Academic Press.

Thrun, S. 2002. Probabilistic robots. *Communications of the ACM* **45**(3) 52–57.

Trummel, K. E., J. R. Weisinger. 1986. The complexity of the optimal searcher path problem. *Operations Research* **34**(2) 324–327.