An Optimal Policy for Target Localization with Application to Electron Microscopy

Raphael Sznitman RAPHAEL.SZNITMAN@EPFL.CH
Ecole Polytechnique Federale de Lausanne, Switzerland

Aurelien Lucchi AURELIEN.LUCCHI@EPFL.CH
Ecole Polytechnique Federale de Lausanne, Switzerland

Peter I. Frazier PF98@CORNELL.EDU
Cornell University, Ithaca, NY 14850, USA

Bruno M. Jedynak BRUNO.JEDYNAK@JHU.EDU
Johns Hopkins University, Baltimore, MD 21218, USA

Pascal Fua PASCAL.FUA@EPFL.CH
Ecole Polytechnique Federale de Lausanne, Switzerland

Abstract

This paper considers the task of finding a target location by making a limited number of sequential observations. Each observation results from evaluating an imperfect classifier of a chosen cost and accuracy on an interval of chosen length and position. Within a Bayesian framework, we study the problem of minimizing an objective that combines the entropy of the posterior distribution with the cost of the questions asked. In this problem, we show that the one-step lookahead policy is Bayes-optimal for any arbitrary time horizon. Moreover, this one-step lookahead policy is easy to compute and implement. We then use this policy in the context of localizing mitochondria in electron microscope images, and experimentally show that significant speed ups in acquisition can be gained, while maintaining near equal image quality at target locations, when compared to current policies.

Proceedings of the 30^{th} International Conference on Machine Learning, Atlanta, Georgia, USA, 2013. JMLR: W&CP volume 28.

1. Introduction

We consider the problem of quickly identifying the location and size of an object of interest by asking a series of questions about the region it occupies. This is key to speeding up the imaging of target objects, such as intracellular structures, when using scanning electron microscopes (SEM), which can focus resources on promising areas (Lucchi et al., 2011; Veeraraghavan et al., 2010). In this application, asking a question involves acquiring image data over a portion of the capture area and feeding it to a classifier that returns an estimate of the presence of a target within the queried region. This data can be either acquired quickly, which degrades the reliability of the estimate, or slowly, which severely impacts overall acquisition time.

While both image acquisition and classification are subject to noise, requiring the use of statistical inference to estimate the target's location, the heart of this work lies in the *Decision Problem* of what questions to ask to localize the target efficiently. As in Active Learning (Dasgupta et al., 2007; Castro et al., 2005; Settles, 2009) and Sequential Experimental Design (DeGroot, 1970; Wetherill & Glazebrook, 1986; Berry & Fristedt, 1985; Srinivas et al., 2010), the decisions can be made adaptively, which may increase efficiency but makes the optimal decision policy more difficult to find.

In this paper, we formulate this problem in a Bayesian

framework and explicitly characterize the globally Bayes-optimal policy. Furthermore, we prove that, somewhat surprisingly, the greedy, or one-step lookahead, policy is in fact Bayes-optimal over arbitrary time horizons. The greedy policy is easy to compute, which is beneficial in many time-sensitive and computationally-limited applications. This optimality result is in contrast with other solutions to Bayesian active learning problems, where the greedy policy is suboptimal and can only serve as a heuristic (Brochu et al., 2009; Zhang et al., 2003).

We demonstrate the effectiveness of this optimal greedy policy in the context of SEM imaging, where our results show that images of desired quality at target locations are acquired in half the time required by other state-of-the-art methods.

Akin to this work, (Jedynak et al., 2011) considers a related target localization task and shows that the greedy policy is also Bayes-optimal over any time horizon. The current work differs from the latter by considering that the target $X^* = (X_1^*, X_2^*)$ covers an interval within the search space, rather than a single point. It also differs in that questions with a variety of noise models and costs are available, while (Jedynak et al., 2011) only considered questions with a single homogeneous cost and noise model. Both aspects are important in the SEM application treated here.

This paper is organized as follows: Sec. 2 surveys the related literature. Then in Sec. 3, we explicitly formulate our problem and objective, and then present our main theoretical result in Sec. 4. In Sec. 5, we detail an embodiment of our policy in the context of SEM imaging, experimentally validating our approach and we conclude with final remarks in Sec. 6.

2. Related Literature

A substantial amount of literature exists in active learning (Dasgupta et al., 2007; Srinivas et al., 2010; Castro et al., 2005; Settles, 2009) and sequential experimental design (DeGroot, 1970; Wetherill & Glazebrook, 1986; Srinivas et al., 2010). Within this broad literature, many authors have considered Bayesian formulations (Brochu et al., 2009; Zhang et al., 2003; Chick & Frazier, 2012; Gittins & Jones, 1974) and it is well-known that the Bayes-optimal policy for a problem in sequential experimental design is the solution to a partially observable Markov decision process (POMDP), which is characterized by the dynamic programming equation. In a few cases authors have applied dynamic programming to explicitly calculate Bayes-optimal sequential policies for specific problems (Berry & Fristedt, 1985; Gittins & Jones, 1974; Chick & Frazier, 2012). In many problems, however, computing the Bayes-optimal sequential policy is considered to be intractable, leading many authors to use greedy policies as heuristics (Brochu et al., 2009; Zhang et al., 2003).

Another line of closely related research is the *Twenty Questions* literature. The original Twenty Questions game and its noise-free answers has long been associated with notions in Information Theory (Cover & Thomas, 1991) and dichotomous search. Later on, a number of works looked at cases where an unknown but bounded number of questions could be answered incorrectly, and where bounds can be computed in some of these cases (Spencer & Winkler, 1992; Dhagat et al., 2004). Related to this work, is that of Distilled Sensing (Haupt et al., 2011), where sparse targets in white Gaussian noise were shown to be reliably found using an adaptive-sampling scheme. More recently, (Jedynak et al., 2011) showed that when all answers are noisy, with a symmetric noise model, choosing questions that include half the probability mass of the posterior distribution yields a globally optimal strategy with regards to the entropy after a finite number of questions. This policy, more generally known as the *probabilistic bisection search* was first proposed by (Horstein, 1963). A number of other works have also analysed its performance, or the performance of closely related policies, under different conditions and measures of uncertainty (Castro & Nowak, 2007; Ben-Or & Hassidim, 2008; Novak, 2008; Chakraborty et al., 2011; Waeber et al., 2011). In this work, we extend in the direction of (Jedynak et al., 2011) and consider the situation where X^* is an interval (as opposed to an unconstrained parameter) and where multiple question types with varying costs can be used to optimize an entropy-based objective.

3. Formulation and Objective

Consider the continuous random vector of dimension two, $X^* = (X_1^*, X_2^*)$. We denote $p_0 = p_0(x_1, x_2)$ for the joint probability density function of X^* and we assume that $p_0(x_1, x_2) = 0$ when $x_1 > x_2$, so that $X_1^* \leq X_2^*$ with probability 1. As depicted in Fig. 1, X^* represents an interval of length $|X_2^* - X_1^*|$. The differential entropy of p_0 (or of X^*) is then the quantity,

$$H(p_0) = -\int_{x_1=0}^{1}\int_{x_2=0}^{1} p_0(x_1, x_2)\log p_0(x_1, x_2)dx_1dx_2 \tag{1}$$

where log is the logarithm base 2. Note that $H(p_0)$ is not necessarily positive.

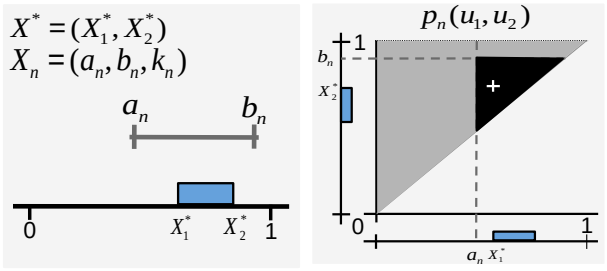

Figure 1. Problem formulation. (*left*) Object to find is an interval, $X^* = (X_1^*, X_2^*)$ (in blue), and is located in (0,1). Example of a possible question $X_n = (a_n, b_n, k_n)$ (in gray) is also shown. (*right*) Probability distribution. Grey region corresponds to the density of X^*, with (X_1^*, X_2^*) shown for this example (white cross). The question $X_n = (a_n, b_n, k_n)$ is also shown with the corresponding $\mathbb{1}_{a \le x_1, x_2 \le b}$ region (in black).

We consider a collection of questions of the form "is $a \le X_1^*$ and $X_2^* \le b$?", indexed by a and b, and where a noisy answer to this question is observed instead of the true one. In addition, we assume that there is in fact a finite collection of such questions, or question *types*, for each pair (a, b). Asking a question of type k, induces a specific cost $W(k)$, and the amount of noise in the observed answers also depends on the type k. Typically, the higher the cost, the lower the noise level.

More formally, a question is specified by a tuple (a, b, k), with $0 < a < b < 1$ and a type $k \in \{1, \dots, K\}$. A given function $k \mapsto W(k) \in \mathbb{R}$ defines the cost of a question of type k. The true answer to a question is denoted $Z(a, b, k)$, where $Z(a, b, k) = 1$ when $a \le X_1^*$ and $X_2^* \le b$, and $Z(a, b, k) = 0$ otherwise. Here, we assume that $Z(a, b, k)$ is unobserved, but that a noisy version, $Y(a, b, k)$ is observed instead. We treat $Y(a, b, k)$ as a continuous random variable whose density given $Z(a, b, k) = 1$ is denoted $f_1^k = f_1^k(y)$ and whose density given $Z(a, b, k) = 0$ is denoted $f_0^k = f_0^k(y)$. Note that the question types are assumed to affect the amount of noise observed in answers, and hence noise models (*i.e.* f_1^k and f_0^k) are indexed by k. In the information theory literature, the mapping $Z(a, b, k) \mapsto Y(a, b, k)$ defines a (noisy) communication channel with binary input alphabet. Such a channel is associated with a *channel capacity* which defines the maximum amount of information that can be transmitted through this channel (Cover & Thomas, 1991). The channel capacity is in this case a scalar in the range $[0; 1]$, where the minimum value (=0) is achieved when $f_0^k = f_1^k$ and the maximal value (=1) is achieved when f_0^k and f_1^k have disjoint support.

We will consider a sequential accumulation of questions and answers. We let $X_n = (a_n, b_n, k_n)$ denote the question asked at time n, and let $Y_n = Y(a_n, b_n, k_n)$ denote the corresponding answer. Starting with an initial distribution on X^* (which is p_0), the first question and answer are $X_0 = (a_0, b_0, k_0)$ and $Y_1 = Y(X_0)$, respectively. The posterior density of X^* given X_0, Y_1 is denoted p_1. Proceeding this way, the history of questions and answers, up to the n^{th} question, is denoted $B_n = \{X_0, Y_1, \dots X_{n-1}, Y_n\}$. We assume that $Y_1, \dots, Y_n$ are conditionally independent given $Z(X_0), \dots, Z(X_{n-1})$, so that if the same question is asked multiple times, potentially different answers are generated due to noise. Then, the posterior density of X^* given B_n is denoted p_n and can be computed recursively using Bayes formulae,

$$p_{n+1}(u_1, u_2) = \frac{p_n(u_1, u_2)}{Q} \left(f_1^{k_n}(Y_{n+1}) \mathbb{1}_{(a_n \le u_1, u_2 \le b_n)} + f_0^{k_n}(Y_{n+1}) \mathbb{1}_{\overline{(a_n \le u_1, u_2 \le b_n)}} \right), \quad (2)$$

where $F_n(a_n, b_n) = p_n(a_n \le X_1^*, X_2^* \le b_n)$, $Q = F_n(a_n, b_n) f_1^{k_n}(Y_{n+1}) + (1 - F_n(a_n, b_n)) f_0^{k_n}(Y_{n+1})$, and the line over the event $a_n \le u_1, u_2 \le b_n$ indicates the complement of that event. From (2), we see that the density p_n is multiplied by a constant factor over the domain $\{(u_1, u_2); a \le u_1, u_2 \le b\}$ and by another constant factor over the complement of this domain, creating a kind of "earthquake" to the surface defined by p_n. The magnitude of this earthquake depends on how different $f_0^{k_n}(y)$ and $f_1^{k_n}(y)$ are from each other.

Objective: In what follows, after having observed the answers to n questions, we are interested in having gained information about X^*. This will be quantified by $H(p_n)$; lower values being preferable. We are also interested in the cost (or time) paid for achieving this information gain. This is quantified by the sum of the costs (or times) of each question; $T_n = \sum_{i=0}^{n-1} W(k_i)$. To cast this problem into the classical dynamic programming setting, we define a value function,

$$V(p, n, t) = \inf_{\pi} E^{\pi} \left[H(p_N) + \lambda T_N | p_n = p, T_n = t \right], \quad (3)$$

where $\lambda \ge 0$ is a constant used to modulate the relation between entropy and cost. A policy which achieves this value is considered optimal. The Optimality Principle (Dynkin & Yushkevich, 1979), allows us to re-write the above value function recursively,

$$V(p, n, t) = \inf_{X_n} E^{Y_{n+1}} \left[V(p_{n+1}, n+1, T_{n+1}) | p_n = p, T_n = t, X_n \right] \quad (4)$$

where the optimization is now only over all possible questions X_n and the expectation is over the answers

that may arise from the question X_n. One great benefit of the recursive definition of the value function is that a policy which attains its minimum, for all n, p, and t, also attains the minimum of (3).

4. Optimal and Greedy Policy

With the goal of solving (3), this section begins by demonstrating how to solve the much simpler, one-step lookahead problem. The derived greedy policy is then shown to be optimal with respect to (3).

First, we define some notation. Let $\varphi(u,k) = H(uf_1^k + (1-u)f_0^k) - uH(f_1^k) - (1-u)H(f_0^k)$, which will be used later to quantify the mutual information, and let $G(u,k) = \varphi(u,k) - \lambda W(k)$, which is called the gain function. Let $G^\star = \max_{u\in[0,1],k\in\{1,\dots,K\}} G(u,k)$. In (Jedynak et al., 2011), it was shown that $u \mapsto \varphi(u,k)$ is concave, and thus this maximum is attained. Let u_k^* be the value of u attaining this maximum given k.

We now state the following lemma, which solves the one-step lookahead problem, and thus implicitly defines the greedy policy.

Lemma 4.1. *For each n,*

$$\min_{X_n} E^{Y_{n+1}}\left[H(p_{n+1}) + \lambda T_{n+1} | X_n, p_n, T_n\right] = H(p_n) + \lambda T_n - G^*. \quad (5)$$

Moreover, any $X_n = (a_n, b_n, k_n)$ attaining the minimum in (5) satisfies $F_n(a_n,b_n) = u_{k_n}^*$ and $k_n = \arg\max_{k\in\{1,\dots,K\}} G(u_k^*, k)$.

Proof. First, $E^{Y_{n+1}}\left[H(p_{n+1}) + \lambda T_{n+1} | X_n, p_n, T_n\right]$ can be written as

$$H(p_n) - I(X^*, Y_{n+1} | X_n, p_n) + \lambda T_n + \lambda W(k_n),$$

where $I(X^*, Y_{n+1}|X_n, p_n)$ is the mutual information between X^* and Y_{n+1} given the history B_n and the question X_n (using the notation of (Cover & Thomas, 1991)). Since $H(p_n)$ and T_n do not depend on X_n, (5) is equivalent to $H(p_n) + \lambda T_n - \sup_{X_n}\{I(X^*, Y_{n+1}|X_n, p_n) - \lambda W(k_n)\}$.

Letting $u = F_n(a_n, b_n)$, we rewrite the mutual information $I(X^*, Y_{n+1}|X_n, p_n)$ as $\varphi(u, k_n)$, and we see that $I(X^*, Y_{n+1}|X_n, p_n) - \lambda W(k_n) = \varphi(u,k_n) - \lambda W(k_n) = G(u, k_n)$. Hence, (5) is equivalent to

$$H(p_n) + \lambda T_n - \sup_{X_n} G(F_n(a_n,b_n), k_n)$$

Since X^* is a continuous random vector, $F_n(a_n,b_n)$ is a continuous function of (a_n, b_n), and there is at least one pair (a_n, b_n) for which $F_n(a_n,b_n) = u_k^*$. Given this value for $F_n(a_n,b_n)$, the optimal value of k_n is $\arg\max_{k\in\{1,\dots,K\}} G(u_k^*, k)$. □

Given the greedy policy defined in the previous lemma, we may now verify that it attains the minimum in (4), for each n.

Theorem 4.2. *Any policy that chooses each $X_n = (a_n, b_n, k_n)$ to satisfy, $F_n(a_n,b_n) = u_k^*$ and $k_n = \arg\max_{k\in\{1,\dots,K\}}\{\varphi(u_k^*, k) - \lambda W(k)\}$ is optimal. Additionally, for each n, the value function is*

$$V(p_n, n, T_n) = H(p_n) + \lambda T_n - (N-n)G^*. \quad (6)$$

Proof. We proceed by backwards induction. At $n = N$ the value function has the claimed form, and the optimal policy makes no decision, and so the claimed form for the optimal policy also holds, vacuously.

For $n < N$, we assume that the value function and optimal policy are of the form claimed at $n+1$. Then,

$$\begin{aligned} V(p_n, n, T_n) &= \inf_{X_n} E^{Y_{n+1}}\left[V(p_{n+1}, n+1, T_{n+1}) | X_n, p_n, T_n\right] \\ &= \inf_{X_n} E^{Y_{n+1}}[H(p_{n+1}) + T_{n+1} - (N-n-1)G^* | X_n, p_n, T_n] \\ &= \inf_{X_n} E^{Y_{n+1}}\left[H(p_{n+1}) + T_{n+1} | X_n, p_n, T_n\right] - (N-n-1)G^* \quad (7) \\ &= H(p_n) + T_n - G^* - (N-n-1)G^* \\ &= H(p_n) + T_n - (N-n)G^*, \end{aligned}$$

which is the claimed form for the value function at time n. In this sequence of equations, we used Lemma 4.1 to show that the infinimum in (7) has the form claimed on the next line. Moreover, the same lemma also shows that the infimum in (7) is attained by an X_n satisfying the form claimed in the statement of the theorem, and thus, by the dynamic programming optimality principle, this X_n is the decision of the optimal policy at time n. Thus, by induction, the claimed form for the value function and optimal policy hold for all n. □

Given this last result, we have obtained an optimal policy which is also a greedy policy and hence easy to implement. It allows one to balance entropy and time for localizing an interval on a number line. Note that a single type of question, which can be computed beforehand, is used at all steps of the optimal policy.

5. Efficient EM Scanning Strategy

In this section, we apply our main result from Sec. 4 to a real-world application. We consider the problem of acquiring as fast as possible SEM images that are

sharp at locations of interest and potentially blurrier elsewhere, and where the speed gains come from not spending imaging resources on uninteresting regions.

For many biologists, SEM offers the ability to acquire detailed images of intra-cellular structures at *nm* resolution, such as those of Fig. 2 (*left*). To this end, this microscope uses an electron beam to sweep the surface of a tissue block, line-by-line, many times. At each pixel location, the average response is computed and used to produce a precise image for the user. The tissue surface is then sliced off the tissue block, and the process is repeated on a newly visible surface. While this process produces image stacks of invaluable worth, doing so is extremely time consuming, *i.e.* 48 hours for a $10\mu m^3$ tissue volume.

Moreover, for many users, imaging the entire block accurately is unnecessary. Instead, they may only be interested in obtaining high resolution images at locations of specific structures of interest. For example, mitochondria are organelles that occupy less than 8%-10% of a tissue sample and have been linked to a number of diseases. Visualizing their development is highly relevant to understanding the progression of illness and recovery. As such, we demonstrate in this section, that images with good resolution at mitochondrion locations can be acquired much faster by our proposed policy when compared to current microscope policies. The remainder of this section details the embodiment of our policy within this context.

5.1. Formulation

We let each slice of a tissue block be composed of a fixed number of *strips*: $\mathcal{S} = \{S_1, \ldots, S_T\}$, where $S_t \in \mathcal{I}$ is a gray scaled image to be imaged (30×960 pixels, Fig. 2 (*left*) depicts these strips). We will apply the policy of Sec. 4 on a strip by strip basis.

On each strip, we are looking for a mitochondrion, denoted by $X^* = (X_1^*, X_2^*) \in S_t$ and which is restricted to be at locations $\{0, 2^d - 1\}$, with $d = 5$, such that each discrete location defines a 30×30 pixel region of the strip (*i.e.* X_1^* is the starting point and X_2^* the end point of the target on one strip). Here, X^* is a discrete random variable with density $p_0(u_1, u_2), u_1, u_2 \in \{0, 2^d - 1\}$ and where $p_0(u_1, u_2) = 0$ for $u_1 > u_2$. The questions are denoted by $X_n = (a_n, b_n, k_n), 0 \leq a_n \leq b_n \leq 2^d - 1$, where (a_n, b_n) are integers, $k_n \in \{1, \ldots, 4\}$ for $0 \leq n \leq N$. The answers Y_{n+1} are computed by first scanning the region interval (a_n, b_n), $W(k_n)$ times. From this, each pixel location of the interval (a_n, b_n) receives $W(k_n)$ values, which are averaged to form a noisy image denoted by I_{a_n, b_n, k_n}. From I_{a_n, b_n, k_n}, we compute a question score, denoted $y_{n+1} \in \mathbb{R}$. In the following subsection we detail how these scores are computed and what their associated noise models are. The cost associated with a question type k_n is as follows: $W(1) = 4, W(2) = 20, W(3) = 40$ and $W(4) = 60$.

Given that X^* is now a discrete, and not a continuous random variable, we are no longer guaranteed to always find an interval question satisfying $F_n(a_n, b_n) = u^*_{k_n}$ and hence the greedy optimal solution from the previous section no longer holds. To approximate the optimal policy in this context, we select the *available* interval and question type that maximizes the gain function $G(u, k)$ instead and compute this by performing a brute force search over the (a, b, k).

5.2. Image Observation and Noisy Channels

Given a sampled image, $I_{a,b,k}$, we must evaluate a question score: y_{n+1}. To do this, we first evaluate if each 30 by 30 pixel block of $I_{a,b,k}$ is part of a mitochondrion or not. This is achieved by first dividing a block into nine 10 by 10 regions, computing an intensity histogram from the center region and constructing another intensity histogram from the combined eight neighbouring regions. These histograms are concatenated and then evaluated with a two-class RBF kernel SVM. Having done so for every block in $I_{a,b,k}$, the proportion of locations where the SVM has returned 1, that is the positive class, is returned.

Given that the number of scans performed, $W(k_n)$, induces significantly different noise levels in $I_{a,b,k}$, a separate SVM must be learned and used for images with $W(k_n)$ scans. Each SVM was trained using five, 1032×1032 pixel images, where an expert provided mitochondria groundtruth segmentations.

Learning observation models: Recall from Sec. 3 that the answer Y_{n+1} comes from one of two possible observation sources: f_1^k or f_0^k. We model both of these to be Gaussian with parameters (μ_1^k, σ_1^k) and (μ_0^k, σ_0^k), respectively and where the parameters are estimated from the same training data used to train the SVMs. To estimate these for any value k, we randomly generated 2000 intervals (a, b) for each image, evaluated $X = (a, b, k)$ and computed the corresponding scores from the 2000 intervals. Using the scores and the mitochondria expert segmentations, we then computed the parameters for f_0 and f_1 using maximum likelihood estimation. Finally, estimating the $\varphi(u, k)$, was achieved by Monte Carlo simulation and we have plotted $G(u, k)$, in Fig. 2 (*right*). Note that in the case where $\lambda = 4.5 \times 10^4$ (which is task specific), $G(\cdot, k)$ can be maximized when $k = 4$ or 20 (*i.e.* using 20

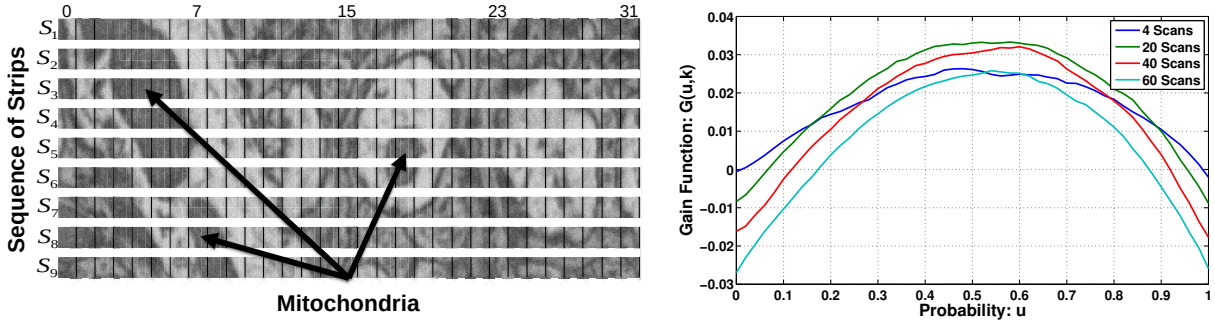

Figure 2. (*left*) Example of an image acquired with a SEM. Our approach localizes the indicated mitochondria one strip, S_i at a time. (*right*) Plot of gain function $G(u,k)$ as a function of the probability u, for different values of $k = \{1,\ldots,4\}$, when $\lambda = 4.5 \times 10^4$.

scans when $0.15 < u < 0.88$ and using 4 scans otherwise) and using other scanning quantities would be sub-optimal.

5.3. Algorithm

The proposed policy can be cast as a *hypothesis-and-test* algorithm. For each strip, we evaluate the algorithm outlined in Alg. 1. By the end of the procedure, we return the posterior distribution, p_N and the set of image observations, $\{I_{a_n,b_n,k_n}\}_{n=0}^N$.

Given that for the first strip S_1, no information on the target location is known, we assign p_0 to be a discrete uniform probability distribution in the range (0,31). This initial prior indicates that the target has equal chance of being anywhere. However, once a strip has been processed with Alg. 1, we maintain the posterior distribution and use it as initial prior for the subsequent strip. That is, the posterior computed after N questions on strip S_t, is the the prior used for strip S_{t+1}. Given that mitochondria are 2D objects (even 3D if the volume is taken into account), and do not move significantly from one strip to another, reusing the computed posteriors from previous strips effectively encodes 2D spatial information.

Note that in practice multiple or no mitochondria may be present on a strip. To deal with this, we use the same approach as in (Sznitman & Jedynak, 2010) and extend the hypothesis space of X^* such that an additional point mass indicates that the target is not present in the visible domain. For multiple targets, we apply non-maximum suppression (on both the imaging domain and the density) to regions that have been observed more than a number of times (*i.e.* 600 scans) and proceed until $n = N$.

Algorithm 1 Fast Strip Imaging (N, p_0, W, λ)

1: **for** $n = 0,\ldots,N$ **do**
2: $\{a_n, b_n, k_n\} = \arg\max_{a,b,k}\{\varphi(F_n(a,b),k) - \lambda W(k)\}$, where $F_n(a,b) = p_n(a \le X_1^*, X_2^* \le b)$
3: Sample noisy image: I_{a_n,b_n,k_n}
4: Compute question score y_{n+1} from I_{a_n,b_n,k_n}
5: Compute posterior distribution p_{n+1} using (2)
6: **end for**
7: **return** p_N and $\{I_{a_n,b_n,k_n}\}_{n=0}^N$

5.4. Experiments

We validate our approach by evaluating our policy against what is typically possible with current microscopes, and a similar policy to that in (Jedynak et al., 2011) which does not take into account time costs. The standard policy scans each location of the block surface a fixed number of times, and averages the received values at each location. For any scanning policy, we may quantify its performance by 1) the average scanning time per pixel and 2) the per-pixel estimated intensity value at mitochondrion locations (measured by the Peak Signal-to-Noise Ratio (PSNR) and the average proportion of error per pixel, *i.e.* the difference between the true and estimated intensity values, divided by the true intensity).

We let the entropy-cost factor $\lambda = 4.5 \times 10^4$ when evaluating each strip and will vary the number of iterations, N. Once a surface imaged, the process is repeated on the next surface.

5.4.1. SEM Data and Image sampling

Using an SEM, we first collected five images of size 1032×1032 pixels, using 600 scans, of a rodent brain tissue. These images were used for training purposes. We then collected 160, 1500×960 images with 600 scans (which when stacked forms a volume) for testing purposes. On both the train and test data, mi-

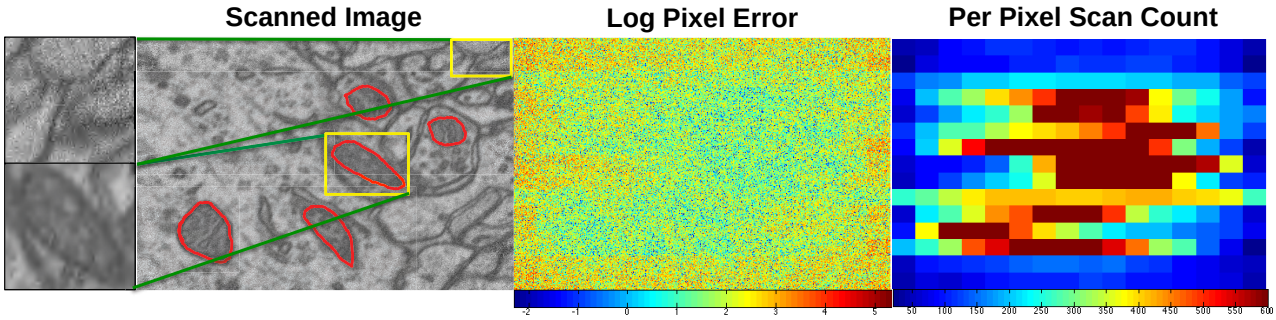

Figure 3. **Scanned Image:** shows the image acquired when running our approach on a series of strips. Red regions indicate mitochondria. Two blown-up regions are shown on the left, where an image regions is noisy (top) and where a large number of scans took place (bottom). **Log Pixel Error:** corresponding image where at each pixel the log pixel difference between our produced image and the true image was computed. **Per Pixel Scan Count:** visually depicts regions that were scanned more times (in orange and red).

tochondria locations were hand segmented by an expert. Given that current microscopes are incapable of storing images acquired with different scan counts for entire volumes, the following scheme (similarly to(Sznitman et al., 2012; Veeraraghavan et al., 2010)) was used to acquire data with $W(k)$ scans at test time [1].

On a single surface, we first collected 60 images acquired with 10 scans each, and note that the corresponding 600 scan image is the per-pixel average of the 60 images, which we denote μ_p for pixel p. The standard deviation of the pixel intensity observed, denoted by σ_p, can be estimated as function of the true pixel intensity, $\sigma_p = m\mu_p + b$, where m and b are linear regression parameters and were estimated from this data. Then, with an image of 600 scans, and assuming a Gaussian model for the pixel intensity values, we can generate an image with $W(k)$ scans by sampling $\mathcal{G}(\mu_p, m\mu_p + b)$, $W(k)$ times, and averaging the received values. This generating process is used to create sample noisy images in Alg. 1 (line 3).

5.4.2. Results

The "Scanned Image" of Fig. 3 shows an example of a produced image using our approach. This image was produced by spending only 280 scans per pixel on average, as opposed to the traditional 600. Mitochondria locations are shown in red, and looking closely at the indicated yellow boxes, which have been blown-up on the left, we notice that the image noise at mitochondria locations is noticeably lower than in other regions. To further show this, the "Log Pixel Error" of Fig. 3 shows the log pixel difference between the true and produced image. Darker regions indicate lower pixel error. Similarly, "Per Pixel Scan Count" of Fig. 3 shows how many times each pixel was scanned. Notice that in both cases the mitochondria locations correlate highly with low pixel error and large scan counts. Submitted with this manuscript, a supplementary video shows the algorithm as it proceeds to scan an entire image, strip by strip. We also show a visual evolution of the posterior distribution after each iteration and the produced image.

We evaluated each approach on the 160 test images, which consists of running our policy 8000 times. Quantitatively, our policy improves on the time necessary to acquire images of equal quality at mitochondria locations, as can be seen from Table. 1. In fact, in nearly half of the time typically taken, we acquire the images of nearly equal quality at mitochondria locations. In practice, this speed up would allow a biologist to scan a $10\mu m^3$ tissue volume and verify a hypothesis in one day instead of two. In addition, we see the benefit of including time cost in our objective function, as we outperform in time and accuracy the case where $\lambda = 0$ and $W(k) = 60$.

Finally, in addition to having high quality images, biologists are often interested in extracting 3D segmentations of mitochondria in order to better visual their 3D structure. One method to do this automatically is to use the MRF method of (Lucchi et al., 2011), and which we have applied to the image stacks produced by our scanning policy. Using the first 80 images for MRF training, we evaluated the remaining 80 images. We report that when imagining with 600 scans, segmentations of 98% accuracy (and a PASCAL VOC score of 83%) are achieved, while our policy, in

[1] Implementing our policy on an SEM is currently impossible largely due to the fact that internal access and details to the device are not available.

Method	Average # of Scans	PSNR	Proportion of per Pixel Error
Full Scanning	600 (0.0)	42.66	0.032 (1.6)
Full Scanning	300 (0.0)	41.10	0.046 (2.3)
Optimal Policy ($N = 80$)	320 (2.7)	**42.50**	**0.034** (2.6)
Optimal Policy ($N = 40$)	**296 (3.1)**	41.89	0.039 (2.7)
Full Scanning	200 (0.0)	40.20	0.056 (2.8)
Optimal Policy ($N = 20$)	**198 (2.5)**	**41.03**	**0.047** (2.5)
Policy ($N = 80, \lambda = 0, W(k) = 60$)	421 (2.6)	42.57	0.034 (2.1)
Policy ($N = 40, \lambda = 0, W(k) = 60$)	380 (2.9)	41.92	0.038 (4.6)
Policy ($N = 20, \lambda = 0, W(k) = 60$)	361 (3.1)	41.39	0.043 (7.1)

Table 1. Quantitative comparison of proposed policy against other polices. The PSNR and Proportion of per Pixel Error (standard error in brackets with a factor of 10^{-6}) are estimated at target locations.

half that time, produces segmentations of 97% accuracy (and a PASCAL VOC score of 81%). Please visit http://cvlab.epfl.ch/~sznitman for additional results and videos.

6. Conclusion

This paper considered the task of finding a target when making a finite number of sequential observation, N, by evaluating intervals of chosen length, position and type. With the constraint that observations are always noisy, and that different types of questions may induce different costs, we have presented a Bayes optimal policy for an entropy-cost based objective function. The use of our optimal policy, which has the particularity of being greedy, is demonstrated in the context of localizing mitochondria in SEM images. We show experimentally that our policy provides significant speed ups in image acquisition while maintaining near equal image quality at target locations, when compared to current policies.

Acknowledgements

This work was funded in part by the ERC MicroNano Grant. Peter Frazier acknowledges support from NSF Award 142251 and AFOSR Award FA9550-11-1-0083.

References

Ben-Or, M. and Hassidim, A. The bayesian learner is optimal for noisy binary search. In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science*, pp. 221–230, 2008.

Berry, D.A. and Fristedt, B. *Bandit Problems: Sequential Allocation of Experiments*. Chapman & Hall, London, 1985.

Brochu, E., Cora, M., and de Freitas, N. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. Technical Report TR-2009-023, Department of Computer Science, University of British Columbia, November 2009.

Castro, R. and Nowak, R. *Active sensing and learning, in Foundations and Applications of Sensor Management*. pringer-Verlag, 2007.

Castro, R., Willett, R., and Nowak, R. Faster rates in regression via active learning. In *NIPS*, 2005.

Chakraborty, M., Das, S., and Magdon-Ismail, M. Near-optimal target learning with stochastic binary signals. In *UAI*, pp. 69–76, 2011.

Chick, S.E. and Frazier, P.I. Sequential sampling for selection with economics of selection procedures. *Management Science*, 58:550–569, 2012.

Cover, T. M. and Thomas, J.A. *Elements of Information Theory*. Wiley Interscience Press, 1991.

Dasgupta, S., Hsu, D., and Monteleoni, C. A general agnostic active learning algorithm. In *NIPS*, 2007.

DeGroot, M. H. *Optimal Statistical Decisions. l, New York*. McGraw Hill, 1970.

Dhagat, A., Gacs, P., and Winkler, P. On playing twenty questions with a liar. Technical report, Boston University, 2004.

Dynkin, E. B. and Yushkevich, A. A. *Controlled Markov Processes*. Springer, 1979.

Gittins, J. C. and Jones, D. M. A dynamic allocation index for the sequential design of experiments. In Gani, J. (ed.), *Progress in Statistics*, pp. 241–266, Amsterdam, 1974. North-Holland.

Haupt, J., Castro, R. M., and Nowak, R. Distilled sensing: Adaptive sampling for sparse detection and estimation. *IEEE Transactions on Information Theory*, 57(9), 2011.

Horstein, M. Sequential decoding using noiseless feedback. *IEEE Transactions on Information Theory*, 9:136–143, 1963.

Jedynak, B., Frazier, P. I., and Sznitman, R. Twenty questions with noise: Bayes optimal policies for entropy loss. *Journal of Applied Probability*, 49:114–136, 2011.

Lucchi, A., Smith, K., Radhakrishna, A., Knott, G., and Fua, P. Supervoxel-Based Segmentation of Mitochondria in Em Image Stacks with Learned Shape Features. *TMI*, 31(2):474–486, 2011.

Novak, R. Generalized binary search. In *Conference on Communication, Control, and Computing*, pp. 568 – 574, 2008.

Settles, B. Active Learning Literature Survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.

Spencer, J. and Winkler, P. Three thresholds for a liar. *Combinatorics, Probability and Computing*, 1:81–93, 1992.

Srinivas, N., Krause, A., Kakade, S., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, 2010.

Sznitman, R. and Jedynak, B. Active Testing for Face Detection and Localization. *PAMI*, 32(10):1914–1920, June 2010.

Sznitman, R., Lucchi, A., Pjescic-Emedji, N., Knott, G., and Fua, P. Efficient Scanning for EM Based Target Localization. In *MICCAI*, 2012.

Veeraraghavan, Ashok, Genkin, Alex, Vitaladevuni, Shiv, Scheffer, Lou, Xu, Shan, Hess, Harald, Fetter, Richard, Cantoni, Marco, Knott, Graham, and Chklovskii, Dmitri. Increasing Depth Resolution of Electron Microscopy of Neural Circuits using Sparse Tomographic Reconstruction. In *CVPR*, pp. 1767–1774, 2010.

Waeber, R., Frazier, P. I., and Henderson, S. G. A bayesian approach to stochastic root finding. In *Winter Simulation Conference*, 2011.

Wetherill, GB and Glazebrook, KD. *Sequential Methods in Statistics*. Monographs on Statistics and Applied Probability. Chapman & Hall, London, third edition, 1986.

Zhang, Yi, Xu, Wei, and Callan, Jamie. Exploration and exploitation in adaptive filtering based on bayesian active learning. In *Proceedings of the Twentieth International Conference (ICML 2003)*, pp. 896–903, Washington, DC, USA, 2003.