Optimal Two-Stage Search for Sparse Targets Using Convex Criteria

Eran Bashan, Raviv Raich, *Member, IEEE*, and Alfred O. Hero, *Fellow, IEEE*

Abstract—We consider the problem of estimating and detecting sparse signals over a large area of an image or other medium. We introduce a novel cost function that captures the tradeoff between allocating energy to signal regions, called regions of interest (ROI), versus exploration of other regions. We show that minimizing our cost guarantees reduction of both the error probability over the unknown ROI and the mean square error (MSE) in estimating the ROI content. Two solutions to the resource allocation problem, subject to a total resource constraint, are derived. Asymptotic analysis shows that the estimated ROI converges to the true ROI. We show that our adaptive sampling method outperforms exhaustive search and are nearly optimal in terms of MSE performance. An illustrative example of our method in radar imaging is given.

Index Terms—Adaptive sampling, adaptive sensing, energy allocation, search methods, sparse signals.

I. INTRODUCTION

THIS paper considers the problem of detecting and estimating signals in an unknown region of interest (ROI), under resource constraints. We formulate this problem as a sequential decision problem: at each iteration, information acquired is used to estimate and refine the ROI.

This problem arises in many applications including: target detection and classification, computer-aided diagnosis, and screening. For example, in a radar reconnaissance mission, active radar may be used to image a given scene. A typical system is designed to detect targets exceeding a minimal profile. This minimal target profile dictates the scan parameters such as the energy level the radar transmits and the scan time duration. Moreover, targets usually occupy a small section of the scanned area, called the ROI. Most systems consider exhaustive search with equal energy allocation, defined as a search policy, where all cells are searched with equal effort allocated to each cell, to spread sensitivity over all locations. As a result, a relatively large portion of the energy is allocated outside the ROI. This excess energy could be used to better illuminate the ROI. Furthermore, in surveillance applications, by deploying energy

over an unknown area, the searcher risks exposure. Reducing the scan energy outside the ROI reduces exposure risks.

As another application, consider the task of early detection of tumors using medical imaging, e.g., X-ray computed tomography. Early detection is concerned with detecting small tumors. In many cases, little *a priori* knowledge about the tumor location exists. Consider the area containing the tumor as an unknown ROI. Lifetime radiation exposure constraints limit the total energy that can be used in a specific scan. There are two questions: a) Where are tumors located? and b) What kind of tumors does the ROI contain? This combined detection/estimation suggests using adaptive sampling over the image to improve both detection and estimation performance.

The search problem considered in this paper bears some similarity to Posner's work on minimizing expected search time for finding a satellite lost in the sky [1]. Posner suggests a two-step procedure. First, briefly search the entire domain to generate a likelihood function of possible satellite locations. Secondly, sequentially search the domain again in order of decreasing likelihood of satellite position. Posner's model assumes that the search is stopped as soon as the satellite has been found and that detection probability is increasing with search time. Therefore, sequentially searching the cells with the highest likelihood reduces the overall expected search time. By minimizing expected search time Posner imposes a "soft" resource constraint on the total time used to search each cell. In this paper, we adopt a Bayesian framework for sequential search for multiple objects obscured by noise. Thus the posterior distribution of objects' presence replaces the likelihood of object presence used in [1]. The use of the posterior in place of the likelihood guarantees minimization of an average cost and a stronger sense of optimality.

Although we search for multiple targets within a signal, we focus on applications where the total support of the signal part containing targets is small compared to the entire signal support. Such signals can be viewed as sparse, and we define the sparsity parameter p as the proportion of the signal support containing targets. Johnstone and Silverman consider a Bayesian framework for estimating sparse signals from a single measurements vector in [2]. They consider thresholding the vector entries and setting the estimated signal equal zero for all measurements below the threshold. Thus, significant gains in estimation mean-square error (MSE) are achieved for small p. They utilize a Bayesian framework to find an optimal threshold minimizing the average squared error. We also use the Bayesian framework to find an optimal threshold. However, we use this threshold to generate additional measurements of all signal elements exceeding it. This adaptive sampling approach and added measurements enables higher gains in reducing estimation MSE.

E. Bashan and A. O. Hero are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: bashan@umich.edu).

R. Raich is with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331 USA (e-mail: raich@eecs.oregonstate.edu).

Wipf and Rao use sparse Bayesian learning in the problem of basis selection from an overcomplete dictionary [3]. They use a parameterized prior on the basis weight coefficients to encourage sparse solution to an otherwise l_2-norm minimization problem. The parameters of the prior are estimated from the data along with the basis weights using an expectation-maximization (EM) algorithm. However, the EM algorithm uses a single measurement of the underlying signal. In this paper, the posterior distribution is used to partially remeasure the underlying signal. By spatially focusing the measurement process onto the ROI, we can better estimate the correct signal values.

Adaptive sampling or active learning for the purpose of estimating functions in noise has been considered in [4]. Castro *et al.* show that for piecewise constant functions, active learning methods can capitalize on the highly localized nature of the boundary by focusing the sampling process in the estimated vicinity of the boundary. Thus, estimation error converges at higher rates than with passive sampling. A two-step approach is proposed where first a rough estimate of the underlying signal is generated and second the signal is resampled in the vicinity of the estimated boundaries. We use a similar two-stage approach where previous measurements are used to determine where to sample next. However, our work differs in two aspects, as we do not limit the discussion to a class of spatially inhomogeneous signals and we consider the additional aspect of resource allocation. While [4] assumes identical sampling procedure for all samples, e.g., similar measurement noise variance, we consider different sampling procedures between stages and among different spatial locations within a particular stage.

Resource allocation in the context of adaptive waveform amplitude design for estimating parameters of an unknown medium under average energy constraints is discussed in [5]. Rangarajan *et al.* derive an optimal amplitude allocation for the second stage in a two-stage problem as a function of the first stage measurements. In this paper, we consider the more general problem of waveform design under a total energy constraint. Therefore, measurements at the first stage are used to select the optimal waveform among all possible waveforms with some bounded total energy. Thus we are able to focus the sampling process onto the ROI.

The cost function considered here is similar in nature to a terminal/total reward utility function. This formulation is used in some multiarm bandit (MAB) problems (see [6, pp. 123]). The difference between MAB and our formulation of the search problem is that each action we take affects the posterior probability distribution of target locations, and our method is greedy.

In this paper, we focus on adaptively determining the ROI that contains targets of interest. Two main contributions in this paper are: 1) we introduce a novel convex cost function for optimizing the ROI search and 2) we provide two-stage optimal and suboptimal adaptive search policies with respect to our cost function. Remarkably, this leads to solutions that minimize both the Chernoff bound on error probability and the *Cramér–Rao* bound on estimating the parameter values within the ROI. The optimal and suboptimal policies are greedy search algorithms with complexity order proportional to the discrete signal support Q. The optimal policy rank orders the posterior distribution values and then finds an optimal threshold and assigns additional effort to
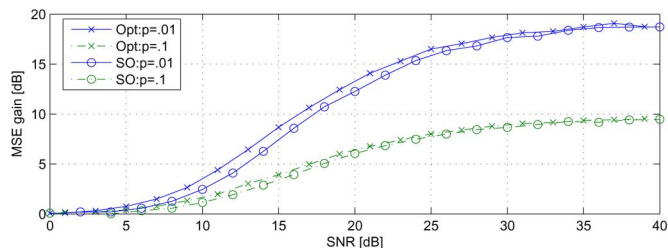

Fig. 1. Gain in MSE for the CME in (46) based on an adaptive search compared to the MSE (variance) of the CME for an exhaustive search policy (45). Curves with crosses correspond to adaptive resource allocation policy, while curves with circles represent the suboptimal adaptive policy, for p values of 1/100 and 1/10. The MSE gain for adaptive resource allocation policy is slightly higher than that of the suboptimal mapping. Note that using our methods results in about 6 dB gain in MSE at SNR value of 13 dB for sparsity level of 1%. In addition, MSE gain is inversely proportional to the sparsity; hence higher gains can be expected for application where $|\Psi| \ll Q$.

all cells with posterior values exceeding the threshold. On the other hand, the suboptimal policy uses the posterior distribution values to assign additional effort to all cells, thus requiring an order $\mathcal{O}(\log Q)$ fewer computations. Both policies outperform an exhaustive search scheme in terms of postprocessing tasks such as target detection and ROI estimation. Moreover, an asymptotic performance analysis at high signal-to-noise ratio (SNR) is given and shows that the estimated ROI converges to the true ROI and the performance gain approaches a theoretical limit, which is inversely proportional to the sparsity of the signal.

Our optimal resource allocation policies are derived for a two-stage resource allocation problem, based on a Gaussian observation model, and assume statistically independent targets. Our methods result in 6 dB performance gain estimating the parameters value within the ROI, at SNR of 13 dB and $p = 1\%$ (see Fig. 1). In addition, the radar imaging example in Section V uses a non-Gaussian model but still results in significant performance gain.

The rest of this paper is organized as follows. Section II formally states the problem and introduces our cost function. In Section III, we present optimal and suboptimal solutions for the resource allocation problem. Section IV includes thorough performance evaluation of the two policies as compared to an exhaustive policy for both detection and estimation. An illustrative example of our methods for a radar imaging system is given in Section V. Finally, we conclude and point to future work in Section VI. For lack of space, highly technical proofs are abbreviated or omitted. They can be found in [7].

II. PROBLEM FORMULATION

Consider a discrete space $\mathcal{X} = \{1, 2, \ldots, Q\}$ containing Q cells and equipped with a probability measure P. We use Ψ to denote an ROI in $\mathcal{X}$, i.e., $\Psi \subseteq \mathcal{X}$. In the sequel, Ψ will be a randomly selected sparse subset of $\mathcal{X}$. We assume $|\Psi| \ll |\Psi^c|$, where $|\Psi|$ equals the number of elements in Ψ and Ψ^c is the relative complement $\mathcal{X} - \Psi$ of Ψ. Exact definition of the ROI is application-dependent. In radar target localization, the ROI is the collection of all cells containing targets and related phenomena, e.g., target shadows. In a medical imaging application, such as early detection of breast cancer, where tumor boundaries

are poorly defined, the ROI may contain all cells containing targets (a tumor) plus some neighboring cells.

Let I_i be an indicator function of the ROI such that

$$I_i = \begin{cases} 1, & i \in \Psi \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

and $\{p_i = \Pr(I_i = 1)\}_{i=1}^Q$ is an associated set of prior probabilities. Let $\boldsymbol{I}_\Psi = [I_1, \ldots, I_Q]'$ be a vector corresponding to the set of all indicators and $(\cdot)'$ denote the transpose operator. We say that the presence of a target affects cell i if $i \in \Psi$. Define the random vector $\boldsymbol{Y} : \mathcal{X} \to \mathbb{R}^Q$ and consider the conditional probability $p(\boldsymbol{Y}|\boldsymbol{I}_\Psi; \lambda)$, where λ is defined below.

Consider a sequential experiment where cell i may be sampled T times. By sampling, we mean that $\boldsymbol{y}(t)$, a realization of $\boldsymbol{Y}$, is observed at time t. Let the distribution $\lambda(i, t) \geq 0$ denote the search effort allocated to cell i at time t, with

$$\sum_{t=1}^T \sum_{i=1}^Q \lambda(i, t) = 1, \quad 0 \leq \lambda(i, t) \leq 1 \tag{2}$$

and $\{\lambda(i, t)\}$ is a mapping from past observations $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{t-1}$ to the probability simplex and is called an effort allocation policy, or, equivalently, a search policy. We focus here on a deterministic mappings λ, although a more general random mapping can also be incorporated into our framework but is not presented here. We assume that the "quality" of the samples is an increasing function of the allocated effort to the associated cell, e.g., measured in terms of Fisher information or inverse variance. In general, effort might be time, computing power, complexity, cost, or energy that is allocated to acquiring a particular cell location. Define the cumulative search effort allotted to cell i as

$$\Lambda(i) = \sum_{t=1}^T \lambda(i, t). \tag{3}$$

Consider the following cost function:

$$J(\Lambda) = \sum_{i=1}^Q I_i A(\Lambda(i)) + \sum_{i=1}^Q B(\Lambda(i)) \tag{4}$$

where $A(\cdot)$ and $B(\cdot)$ are decreasing functions that may depend on additional parameters. This restriction ensures that allocating more effort to cells reduces the overall cost. Note that our cost function (4) depends directly on the ROI via the summand of the first sum on the right-hand side (RHS) of (4). Choosing $B(\cdot) = 0$ emphasizes focusing efforts on the ROI. $A(x) = (2\nu - 1)/x$ and $B(x) = (1 - \nu)/x$, with $\nu \in [1/2, 1]$, simplifies $J(\Lambda)$ to

$$J(\Lambda) = \sum_{i=1}^Q \frac{\nu I_i + (1 - \nu)(1 - I_i)}{\Lambda(i)} \tag{5}$$

which has some intuitive and appealing properties. Minimizing this cost function concentrates ν of the total effort over the ROI Ψ and $(1-\nu)$ to its complement Ψ^c, with ν being the dividing factor. Setting $\nu = 1$ focuses all the effort at the ROI, while $\nu = 1/2$ results in an exhaustive resource allocation policy, i.e.,

equal effort allocated to all cells inside and outside the ROI. Furthermore, the choice $\nu = 1$ connects (5) to known criteria such as the Cramér–Rao and Chernoff bounds. For example, in the context of estimating a deterministic signal $\boldsymbol{\mu}$ in additive Gaussian noise, minimizing (5) is equivalent to minimizing the Cramér–Rao lower bound on $\mathrm{E}\left\{\sum_i I_i(\hat{\mu}_i - \mu_i)^2\right\}$; see [7] for details. In a sense, ν controls the tradeoff between exploitation of signal in the ROI ($\nu = 1$) and exploration of the entire signal ($\nu = 1/2$).

In addition, consider a binary Gaussian hypothesis testing problem. Define the null hypothesis $H_0(i)$ as $\mu_i = 0$ and the alternative $H_1(i)$ as $\mu_i > 0$ with a prior probability $\pi = \Pr(H_1)$. Consider the task of deciding between the two hypotheses $\forall i \in \Psi$. The probability of error, i.e., making the wrong decision, defined as $P_e = \Pr(\text{decide } H_0|H_1)\pi + \Pr(\text{decide } H_1|H_0)(1 - \pi)$, can be broken into two parts: misdetect probability P_m over Ψ and false alarm probability P_{fa} over Ψ^c. With $\nu = 1$, we show in [7] that minimizing (5) is equivalent to uniformly minimizing the Chernoff bound on the probability of error P_m over the ROI. Setting $\nu = 1/2$, most of the energy is spread over Ψ^c due to the assumed sparsity $|\Psi| \ll |\Psi^c|$. This leads to a lower Chernoff bound on P_{fa} or, correspondingly, fewer false alarms. If $\nu \in (1/2, 1)$, we trade the two cases, either relaxing the upper bound on P_m or on P_{fa}. In Section IV, we corroborate this intuition with simulation results, where adaptive measurement policies derived from minimization of (5) are used to generate data that is used for both estimation and detection tasks.

Next, we provide an achievable lower bound on our cost function (5).

Lemma 1: The cost function (5) is lower bounded by

$$J(\Lambda) \geq \left[\sqrt{\nu}|\Psi| + \sqrt{1 - \nu}(Q - |\Psi|)\right]^2. \tag{6}$$

This lower bound is achievable with a two-level effort allocation scheme Λ_o, equal over the ROI and equal over its complement, defined as

$$\Lambda_o(i) = \begin{cases} \frac{\sqrt{\nu}}{\sqrt{\nu}|\Psi| + \sqrt{1-\nu}(Q-|\Psi|)}, & i \in \Psi \\ \frac{\sqrt{1-\nu}}{\sqrt{\nu}|\Psi| + \sqrt{1-\nu}(Q-|\Psi|)}, & i \in \Psi^c. \end{cases} \tag{7}$$

Proof: For a nonnegative series $\{a_i\}$, Cauchy–Schwarz inequality provides

$$\left(\sum_{i=1}^Q \frac{a_i}{\Lambda(i)}\right)\left(\sum_{i=1}^Q \Lambda(i)\right) \geq \left(\sum_{i=1}^Q \sqrt{\frac{a_i}{\Lambda(i)}}\sqrt{\Lambda(i)}\right)^2. \tag{8}$$

Since $\sum_{i=1}^Q \Lambda(i) = 1$, (8) simplifies to

$$\sum_{i=1}^Q \frac{a_i}{\Lambda(i)} \geq \left(\sum_{i=1}^Q \sqrt{a_i}\right)^2. \tag{9}$$

Substituting $a_i = \nu I_i + (1 - \nu)(1 - I_i)$, i.e., $\forall i \in \Psi : a_i = \nu$ and $\forall i \in \Psi^c : a_i = 1 - \nu$, into the RHS of (9) yields

$$\left(\sum_{i=1}^Q \sqrt{a_i}\right)^2 = \left[\sqrt{\nu}|\Psi| + \sqrt{1 - \nu}(Q - |\Psi|)\right]^2. \tag{10}$$

Noting that on the left-hand side (LHS) of (9) we have $J(\Lambda) = \sum_{i=1}^{Q} a_i/\Lambda(i)$ proves the validity of (6). To prove the second part of the lemma, note that

$$J(\Lambda_o) = \sum_{i \in \Psi} \frac{\nu}{\Lambda_o(i)} + \sum_{i \in \Psi^c} \frac{1-\nu}{\Lambda_o(i)} \qquad (11)$$

and algebra yields

$$J(\Lambda_o) = \left[\sqrt{\nu}|\Psi| + \sqrt{1-\nu}(Q - |\Psi|) \right]^2 \qquad (12)$$

which is exactly the RHS of (6). This completes the proof. ∎

1) Discussion: We would like to point out the potential performance gains using our cost function (5). Let Λ_U denote an exhaustive search policy with $\Lambda_U(i) = 1/Q$. If $\nu = 1/2$, then (7) results in $\Lambda_U = \Lambda_o$, i.e., Λ_U is optimal. For a general ν, the cost (5) associated with Λ_U is

$$J(\Lambda_U) = Q\left[\nu|\Psi| + (1-\nu)(Q - |\Psi|)\right] \qquad (13)$$

with the two special cases of

$$J(\Lambda_U) = \begin{cases} \frac{Q^2}{2}, & \nu = 0.5 \\ Q|\Psi|, & \nu = 1. \end{cases} \qquad (14)$$

Since Λ_U does not offer any adaptivity, any good adaptive resource allocation policy Λ should result in $J(\Lambda) \leq J(\Lambda_U)$. Therefore, we define the performance gain for (5) in decibels as

$$G(\Lambda) = -10 \log \frac{J(\Lambda)}{J(\Lambda_U)}. \qquad (15)$$

For $\nu = 1$, $J(\Lambda_o) = |\Psi|^2$ and the optimal gain $G(\Lambda_o) = -10 \log |\Psi|/Q$ is achieved by the omniscient resource allocation policy that knows the ROI and consequently concentrates all efforts in the ROI. Define $p^* = |\Psi|/Q$; then $p^* \to 0$ forces $G(\Lambda_o) \to \infty$. Consequently, a good sampling method should yield large gains in a sparse setting, i.e., when p^* is small, the ROI is small. In the following, we will develop a sampling method that exploits these gains. Taking the derivative of the RHS in (6) with respect to ν, it can be shown [7] that $J(\cdot) \geq |\Psi|^2$ for all $\nu \in [1/2, 1]$ and $|\Psi|/Q \leq 1/2$. In other words, if $|\Psi|/Q \leq 1/2$, the optimal gain is achieved by Λ_o for $\nu = 1$. Unfortunately, the omniscient search policy Λ_o is not feasible since the ROI location is a priori unknown. Our goal is to derive a feasible policy with performance gain approaching $G(\Lambda_o)$ for unknown ROI.

In this paper, we restrict our attention to minimizing the expected value of (5). This probabilistic setting utilizes the conditional distribution $p(\mathbf{Y}|\mathbf{I}_\Psi; \lambda)$ in our model. Assuming we observe realizations of $\mathbf{Y}$, our goal is to find a search policy

$$\hat{\lambda}(i,t) = \arg\min_{\lambda(i,t)} \mathrm{E}\left\{ \sum_{i=1}^{Q} \frac{\nu I_i + (1-\nu)(1-I_i)}{\Lambda(i)} \right\} \qquad (16)$$

where $\Lambda(i)$ is given by (3). Next, we introduce a total energy constraint.

A. Energy Allocation With Energy Constraint

Consider a fixed-target radar measurements in the presence of noise. We assume that a radar transmits energy in some known beam pattern to probe a collection of cells. We further assume that the radar is subject to an energy constraint and that observations obey the following model:

$$y_j(t) = \sum_{i=1}^{Q} h_{ij}(t)\sqrt{\lambda(i,t)}\theta_i(t)I_i + n_j(t), \quad t = 1, 2, \ldots, T \qquad (17)$$

where $h_{ij}(t)$ are known weights corresponding to the beam pattern, $\lambda(i,t)$ is the energy allocated for measuring cell i, $\theta_i(t)$ is a random return from cell i, and $n_j(t)$ is an additive observation noise, all at time t. Note that since the indicator of the ROI I_i is independent of t, this model corresponds to a static target scenario. We assume that the additive noise $n_j(t)$ is independent for different j and t. Also, assume that the positive $\theta_i(t)$ follow a prior distribution $p_\Theta(\theta)$ and are independent for different i but may be dependent for different t. The model in (17) can be written in vector form

$$\mathbf{y}(t) = H(t)'\mathrm{diag}\{\sqrt{\boldsymbol{\lambda}(t)}\}\mathrm{diag}\{\boldsymbol{\theta}(t)\}\mathbf{I}_\Psi + \mathbf{n}(t) \qquad (18)$$

where $\mathbf{y}(t) = [y_1(t), y_2(t), \ldots, y_r(t)]'$, $[H(t)]_{ij} = h(t)_{ij}$, $\boldsymbol{\lambda}(t) = [\lambda(1,t), \lambda(2,t), \ldots, \lambda(Q,t)]'$, $\boldsymbol{\theta}(t) = [\theta_1(t), \theta_2(t), \ldots, \theta_Q(t)]'$, $\mathbf{I}_\Psi = [I_1, I_2, \ldots, I_Q]'$, and $\mathbf{n}(t) = [n_1(t), n_2(t), \ldots, n_r(t)]'$. The notation $\sqrt{\boldsymbol{\lambda}(t)}$ denotes a $Q \times 1$ vector with $[\sqrt{\boldsymbol{\lambda}(t)}]_i = \sqrt{\lambda(i,t)}$, while the operator $\mathrm{diag}\{\mathbf{x}\}$ corresponds to a $Q \times Q$ (square) diagonal matrix with $x(i)$ as its iith element. Qualitatively, our objective is to specify a resource allocation policy $\lambda(i,t)$ that improves the "quality" of the measurements, where λ is subject to a total energy constraint

$$\frac{1}{\lambda_T} \sum_{i=1}^{Q} \sum_{t=1}^{T} \lambda(i,t) = 1. \qquad (19)$$

At time t, the energy allocated to cell i may depend on past observations, i.e., $\lambda(i,t)$ is a function of $\mathbf{y}(1)$, $\mathbf{y}(2)$, $\ldots, \mathbf{y}(t-1)$. For brevity, we use the notation $\lambda(i,t)$ to denote $\lambda(i,t; \mathbf{y}(1), \mathbf{y}(2), \ldots, \mathbf{y}(t-1))$. Following (3), define Λ the cumulative energy distributed to cell i as $\Lambda(i) = 1/\lambda_T \sum_{t=1}^{T} \lambda(i,t)$. Our cost function is $J(\Lambda)$, defined in (5), and our goal is to minimize the expected cost in the RHS of (16) over all possible energy allocations $\lambda(i,t)$, subject to (19). Consider $\nu = 1$ and let $p = \mathrm{Pr}(I_i = 1)$ be a uniform prior distribution on the location of targets, where p represents the sparsity of the vector $\mathbf{I}_\Psi$, i.e., $|\Psi|$ is a binomial random variable (RV) with $\mathrm{E}\{|\Psi|\} = pQ$. Define $p^* = |\Psi|/Q$; then $\mathrm{E}\{p^*\} = p$ and $\mathrm{var}(p^*) = p(1-p)/Q$. To find an upper bound on possible performance gains, we use $J(\Lambda_o) \leq J(\Lambda)$, for any Λ, and thus

$$G(\Lambda) \leq G(\Lambda_o) = -10 \log p^*. \qquad (20)$$

In [7], we use (20) and Bernstein's inequality to prove, for $\delta \geq \exp\{-Q3p/(8(1-p))\}$, that

$$\mathrm{Pr}(G(\Lambda) \leq -10 \log p + \epsilon(\delta, p, Q)) \geq 1 - \delta. \qquad (21)$$

In the radar imaging example from Section V, this yields $\mathrm{Pr}(G(\Lambda) \leq 20 + 0.333) \geq 0.999$.

III. SEARCH POLICY

In the following section, we solve the optimization problem for $T = 2$. The idea is to expend a portion of the total energy at $t = 1$ to learn about the data domain and obtain a posterior distribution. Then use it at stage $t = 2$ to refine the posterior and estimate the ROI. We solve

$$\hat{\lambda}(i,t) = \arg\min_{\lambda(i,t)} \mathrm{E}\left\{ \sum_{i=1}^{Q} \frac{\nu I_i + (1-\nu)(1-I_i)}{\sum_{t=1}^{2} \lambda(i,t)} \right\} \quad (22)$$

subject to a total energy constraint $1/\lambda_T \sum_i \sum_t \lambda(i,t) = 1$. Initially, the prior distribution on targets location is uniform, $p_i = p$ for all i. This could be modified if there were prior knowledge on targets location. Let $\lambda^{(t)}$ be the total energy spent at search stage t with $1/\lambda_T \sum_t \lambda^{(t)} = 1$, $t = 1, 2$. Our goal is to optimize the energy distribution between the two stages and among cells in each stage. See [7] for a discussion of the case of $T > 2$.

A. Optimal Two-Stage Search Policy

With $T = 1$ and a uniform prior, we show in the Appendix that the minimizer of the cost (22) is an equal energy allocation

$$\lambda(i,1) = \frac{\lambda_T}{Q} = \frac{\lambda^{(1)}}{Q} \triangleq \lambda_1. \quad (23)$$

Let Γ be the set of all search policies with $\lambda(i,1) = \lambda_1$. For $T = 2$, we find the optimal search policy $\lambda \in \Gamma$ minimizing (22). Since $\lambda_T = Q\lambda_1 + \lambda^{(2)}$, optimizing the total energy allocated for each stage is equivalent to finding an optimal pair $(Q\lambda_1, \lambda_T - Q\lambda_1)$, which involves minimizing over a single variable. Hence the cost function (22) simplifies to

$$J(\Lambda) = \sum_{i=1}^{Q} \mathrm{E}\left\{ \frac{\nu I_i + (1-\nu)(1-I_i)}{\lambda_1 + \lambda(i,2)} \right\} \quad (24)$$

where expectation is taken with respect to $\boldsymbol{y}(1)$ and $\boldsymbol{I}_\Psi$. Note that $\lambda(i,2)$ does not depend on $\boldsymbol{y}(2)$; thus $\boldsymbol{y}(2)$ can be omitted from the expectation in (24). In addition, (24) is constant in $\boldsymbol{\theta}$, and therefore we omit it from the expectation as well. Rewriting (24) using iterated expectation yields

$$J(\Lambda) = \sum_{i=1}^{Q} \mathrm{E}\left\{ \mathrm{E}\left\{ \frac{\nu I_i + (1-\nu)(1-I_i)}{\lambda_1 + \lambda(i,2)} \middle| \boldsymbol{y}(1) \right\} \right\}. \quad (25)$$

Note that I_i is a binary RV. In addition, given $\boldsymbol{y}(1)$, $\lambda(i,2)$ is deterministic. Hence (25) becomes (26) and (27) as shown at the bottom of the page. Using Bayes rule, we obtain

$$\Pr(I_i = 1|\boldsymbol{y}(1)) = \frac{P(\boldsymbol{y}(1)|I_i = 1)p}{P(\boldsymbol{y}(1)|I_i = 1)p + P(\boldsymbol{y}(1)|I_i = 0)(1-p)}$$
$$\triangleq p_{I_i|\boldsymbol{y}(1)} \quad (28)$$

where $P(\boldsymbol{y}|I_i) = \int P(\boldsymbol{y}|I_i, \theta_i) p_\Theta(\theta) d\theta$ is the given conditional probability model describing the measurement dependency on the target. Finally, we rewrite our cost function on the RHS of (24) and solve

$$\hat{\lambda}(i,t) = \arg\min_{\lambda_1, \lambda(i,2)} \mathrm{E}\left\{ \sum_{i=1}^{Q} \frac{\nu p_{I_i|\boldsymbol{y}(1)} + (1-\nu)(1-p_{I_i|\boldsymbol{y}(1)})}{\lambda_1 + \lambda(i,2)} \right\}. \quad (29)$$

Let λ_1 be the energy allocated to each cell at the first step, with $Q\lambda_1 \leq \lambda_T$. Define the RV

$$W_j = \nu p_{I_j|\boldsymbol{Y}} + (1-\nu)(1 - p_{I_j|\boldsymbol{Y}}), \quad j = 1, 2, \ldots, Q \quad (30)$$

with vector of corresponding realizations $\boldsymbol{w} = [w_1, w_2, \ldots, w_Q]'$. Let $\tau : \mathcal{X} \to X$ be a permutation operator that corresponds to the rank ordering of the w_is

$$\tau(j) = \arg\min_{i=1,\ldots,Q}\{w_i : w_i \geq w_{\tau(j-1)}\}, \quad j \in \{1, 2, \ldots, Q\} \quad (31)$$

with $w_{\tau(0)} \triangleq 0$. Whenever the RHS of (31) is not unique, we select an arbitrary i satisfying $w_{\tau(1)} \leq w_{\tau(2)} \leq \cdots \leq w_{\tau(Q)}$. The solution $\lambda_1, \lambda(i,2)$ of (29) will depend on an integer threshold k_0 that is equal to the number of cells that are not searched at $t = 2$. Assuming $w_{\tau(1)} > 0$, define k_0, the threshold parameter, as $k_0 = 0$ if

$$\frac{\lambda_T}{\lambda_1} > \frac{\sum_{i=1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(1)}}} \quad (32)$$

otherwise $k_0 \in \{1, \ldots, Q-1\}$ is the integer, satisfying

$$\frac{\sum_{i=k_0+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k_0+1)}}} < \frac{\lambda_T}{\lambda_1} - k_0 \leq \frac{\sum_{i=k_0+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k_0)}}}. \quad (33)$$

An optimal solution for the second step $\lambda(i,2)$ is given in the Appendix, and the special case where $w_{\tau(i)} = 0$, for all i, is

$$\sum_{i=1}^{Q} \mathrm{E}\left\{ \frac{\mathrm{E}\{\nu I_i + (1-\nu)(1-I_i)|\boldsymbol{y}(1)\}}{\lambda_1 + \lambda(i,2)} \right\} \quad (26)$$

$$= \sum_{i=1}^{Q} \mathrm{E}\left\{ \frac{\nu\Pr(I_i = 1|\boldsymbol{y}(1)) + (1-\nu)(1 - \Pr(I_i = 1|\boldsymbol{y}(1)))}{\lambda_1 + \lambda(i,2)} \right\} \quad (27)$$

treated there. Its properties are summarized below. Given λ_1 and k_0, it is shown that the optimal energy allocation $\lambda(\tau(i), 2)$ minimizing the cost (29) is of the form

$$
\lambda(\tau(i), 2; \boldsymbol{y}(1), \lambda_1)
$$
$$
= \left(\frac{\lambda_T - k_0 \lambda_1}{\sum\limits_{j=k_0+1}^{Q} \sqrt{w_{\tau(j)}}} \sqrt{w_{\tau(i)}} - \lambda_1 \right) I(i > k_0). \quad (34)
$$

Substituting (34) into (29), the optimization problem is equivalent to finding λ_1^* minimizing (35) and (36) as shown at the bottom of the page, where if $k_0 = 0$, then the first summation in (36) equals zero. We can find λ_1^* via a one-dimensional search. In summary, the optimal policy minimizing (29) is as follows.

Algorithm 1: Two-stage adaptive resource allocation policy (ARAP) λ_A

Step 1: Allocate $\lambda_A(i, 1) = \lambda_1^*$ to each cell and measure $\boldsymbol{y}(1)$.

Step 2: Given $\boldsymbol{y}(1)$ compute posteriors $p_{I_i | y_i(1)}$ defined in (28), then sort the W_is defined in (30).

Step 3: Use λ_1^* and the ordered statistic $w_{\tau(i)}$ to find k_0 using (32) and (33).

Step 4: Given k_0, apply $\lambda(i, 2)$, the energy allocation to cell i, as

$$
\lambda_A(\tau(i), 2) = \lambda(\tau(i), 2)
$$
$$
= \left(\frac{\lambda_T - k_0 \lambda_1^*}{\sum\limits_{j=k_0+1}^{Q} \sqrt{w_{\tau(j)}}} \sqrt{w_{\tau(i)}} - \lambda_1^* \right)
$$
$$
\times I(i > k_0), \quad (37)
$$

and measure $\boldsymbol{y}(2)$.

Note that ARAP is a water-filling algorithm. This is a direct consequence of the fact that the cost function minimized in (29) is convex, $\lambda(i, t)$ is positive, and our constraint is linear of the form $\sum_{i,t} \lambda(i, t) = C$ (see [8, pp. 245]).

B. Properties of the Optimal Energy Allocation

Theorem 1: For some $\nu \in [1/2, 1]$ and λ_1^*, let $\Lambda_A(i)$ be the search policy obtained using ARAP. Then, for a uniform prior distribution, we have

$$
J(\Lambda_A) \leq J(\Lambda_U) \quad (38)
$$

with equality achieved if $P_{I_i | \boldsymbol{y}(1)} = c, \forall i$.

Proof: Note that an equal effort allocation policy can be broken into any arbitrarily number T of consecutive search steps, as long as, for all i

$$
\Lambda_U(i) = \frac{1}{\lambda_T} \sum_{t=1}^{T} \lambda(i, t) = \frac{1}{Q}. \quad (39)
$$

Without loss of generality, let $T = 2$ and let Γ denote the family of all effort allocation policies with $\lambda(i, 1) = \lambda_1$ (equal effort allocation) at the first stage, as shown in (40) at the bottom of the page. Since ARAP yields the optimal effort allocation for any set of posterior distribution $\{P_{I_i | \boldsymbol{y}(1)}\}_{i=1}^{Q}$, we have $J(\Lambda_A) \leq J(\Lambda_U)$. If $P_{I_i | \boldsymbol{y}(1)} = c$ for all i, then $w_i = c', \forall i$, for which ARAP yields $k_0 = 0$. Thus, from (37), we obtain $\lambda_A(i, 2) = \lambda_T / Q - \lambda_1$ or, equivalently, $\Lambda_A(i) = \lambda_T / Q = \Lambda_U(i)$. This completes the proof. ∎

ARAP is optimal over all policies that allocate energy uniformly at the first step. In [7], we give the optimal (ARAP) strategy for a general case of a nonuniform prior. However, in this general case, $\lambda^*(i, 1)$ depends on the specific prior and is a function of i. Therefore, the optimization problem involves searching a $(Q + 1)$-dimensional space and is computationally exhaustive.

1) Asymptotic (High SNR) Properties of ARAP: The following properties are proved in [7], when $\nu = 1$. By asymptotic conditions we mean in the limit of high SNR. We define SNR as the signal-to-noise ratio per cell for an equal energy allocation, i.e., SNR $= (\lambda_T / Q) / \sigma^2$. We show the following:

$$
\lambda_1^* = \arg \min_{\lambda_1} \mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{W_{\tau(i)}}{\lambda_1 + \lambda(\tau(i), 2; \boldsymbol{y}(1), \lambda_1)} \right\} \quad (35)
$$
$$
= \arg \min_{\lambda_1} \mathrm{E} \left\{ \frac{1}{\lambda_1} \sum_{i=1}^{k_0} W_{\tau(i)} + \frac{1}{\lambda_T - k_0 \lambda_1} \sum_{i=k_0+1}^{Q} \sum_{j=k_0+1}^{Q} \sqrt{W_{\tau(i)} W_{\tau(j)}} \right\} \quad (36)
$$

$$
\Gamma = \left\{ \lambda(i, t) : \sum_{i=1}^{Q} \sum_{t=1}^{2} \lambda(i, t) = \lambda_T, \ \lambda(i, 1) = \lambda_1 \geq 0, \ \lambda(i, 2) \geq 0 \right\} \quad (40)
$$

i) the average energy allocation threshold goes to a limit $\mathrm{E}\{k_0\} \rightarrow (1-p)Q$;

ii) the per cell energy assignment at stage one goes to zero $\lambda_1^* \rightarrow 0$;

iii) the performance gain goes to the limit $G(\Lambda_A) \rightarrow 10 \log 1/p^*$.

These asymptotic properties are proved under the asymptotic consistency condition

$$p_{I_i|\boldsymbol{y}(1)} \rightarrow I_i \tag{41}$$

in probability. In [7], we prove the validity of (41) for the Gaussian case, where $y_i(1) \sim \mathcal{N}(\sqrt{\lambda_1}\theta_i I_i, \sigma^2)$, and we speculate that (41) holds for other cases as well. Let $\widehat{\Psi} = \{i : \lambda(i,2) > 0\}$ denote the second stage allocation set. Note that $\nu = 1$ provides $w_i = p_{I_i|\boldsymbol{y}(1)}$ and thus together (41) and i) suggest that $\widehat{\Psi}$ converges to Ψ in the sense that $\lim_{\mathrm{SNR}\to\infty} |\Psi \cap \widehat{\Psi}|/|\Psi| = 1$. In addition, ii) implies that the energy allocated by ARAP to the first stage $\lambda^{(1)} = \lambda_1^* Q \rightarrow 0$, therefore increasing the available energy for stage two. Thus, asymptotically ARAP achieves the gain of the omniscient allocation policy Λ_o, and therefore Λ_A is asymptotically optimal. For large Q, we have $p^* \rightarrow p$ in probability. Hence, in this Q, the asymptotic gain is proportional to $\log 1/p$ and we conclude that higher gains are achieved for sparser signals.

C. Suboptimal Two-Stage Search Policy

Note that ARAP requires sorting the w_is and solving Q inequalities to find the threshold k_o. These operations require an order $\mathcal{O}(\log Q)$ computations. As a simple alternative to ARAP, we consider a search policy where $\lambda(i,1) = \lambda_1$ and

$$\lambda(i,2) = \frac{\lambda_T - Q\lambda_1}{\sum_{j=1}^{Q} \sqrt{w_j}} \sqrt{w_i}$$

leading to a corresponding cumulative energy allocation policy Λ. Substituting Λ in (29) yields a single variable optimization problem requiring simple search over $\lambda_1 \in (0, \lambda_T/Q)$ to find λ_1^* minimizing the expected cost (24). Finally, we define $\lambda_{so}(i,1) = \lambda_1^*$ and

$$\lambda_{so}(i,2) = \frac{\lambda_T - Q\lambda_1^*}{\sum_{j=1}^{Q} \sqrt{w_j}} \sqrt{w_i} \tag{42}$$

with its equivalent cumulative energy allocation Λ_{so}. The simple allocation policy (42) is optimal, i.e., minimizing (29) when $\nu = 1$, for two extreme cases: i) uniform posterior distribution and ii) posterior distribution vector $[p_{I_1|\boldsymbol{y}(1)}, \ldots, p_{I_Q|\boldsymbol{y}(1)}]'$ with L elements for which $p_{I_i|\boldsymbol{y}(1)} = 1$ and $(Q - L)$ elements for which $p_{I_i|\boldsymbol{y}(1)} = 0$. For i), we get an equal energy allocation, while for ii), (42) reduces to

$$\lambda(i,2) = \begin{cases} \frac{\lambda_T - Q\lambda_1}{L}, & p_{I_i|\boldsymbol{y}(1)} = 1 \\ 0, & \text{otherwise} \end{cases} \tag{43}$$

both equivalent to the optimal mapping (37). Although (42) does not make the analytical evaluation of the expectation in (24) tractable, it is less computational demanding than the optimal solution. In fact, since λ_1^* is a function of SNR and the sparsity

p, it can be computed offline. Thus, λ_{so} is a direct mapping from the observation space to the search space. Next, we compare the two policies and show that λ_{so} is nearly optimal in terms of (24).

IV. COMPARISON OF SEARCH ALGORITHMS

Assume either λ_A or λ_{so} were used to generate data vectors $(\boldsymbol{y}(1), \boldsymbol{y}(2))$. A natural question is whether or not this adaptive data acquisition is better in some sense than the nonadaptive acquisition obtained using the standard exhaustive search policy with equal energy constraint. In this section, we compare performance of both the optimal and suboptimal effort allocation policies to those achieved by exhaustive search. We start by showing performance gains in both estimation and detection due to our adaptive measuring schemes. Next, we compare the performance (24) achieved by λ_A and λ_{so} to show that λ_{so} is nearly optimal. In Section IV, we assume $\theta_i(t) = \theta_i$ are independent and identically distributed (i.i.d.) truncated Gaussian random variables with mean $\mu_\theta = 1$ and standard deviation $\sigma_\theta = 0.25$, for all i.

A. Estimation Postprocessing

Consider the problem of estimating the true value of a target return θ_i from $\boldsymbol{y}(t)$ given by

$$y_i(t) = \sqrt{\lambda(i,t)}\theta_i I_i + n_i(t) \tag{44}$$

where $n_i(t) \sim \mathcal{N}(0, \sigma^2)$, $t = 0$ represent an exhaustive search with $\lambda(i,0) = \lambda_T/Q \triangleq \lambda_0$ and $t = 1, 2$ are due to either the optimal or suboptimal measurement policies λ_A or λ_{so}. For estimation, the choice of $\nu = 1$ seems natural. Recall that in Section II, we claimed that minimizing our cost is entirely equivalent, in the Gaussian case, to minimizing the estimation MSE. Assuming $\theta_i \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$, we use a Bayesian framework for estimating θ_i based on its prior distribution. The optimal estimator minimizing the MSE is the conditional mean estimator (CME). We compare the performance of the CME $\mathrm{E}\{\theta_i|\boldsymbol{y}(0)\}$ for an exhaustive search policy to the CME $\mathrm{E}\{\theta_i|\boldsymbol{y}(1), y_i(2)\}$ for either ARAP or the suboptimal search policy. The MSE of the CME for the exhaustive search policy is given by

$$\mathrm{var}\{\theta_i|y_i(0)\} = \sigma_\theta^2 - \frac{\lambda_0 \sigma_\theta^4}{\sigma^2 + \lambda_0 \sigma_\theta^2} = \frac{\sigma_\theta^2}{1 + \lambda_0 \frac{\sigma_\theta^2}{\sigma^2}}. \tag{45}$$

The competing estimator is a naive Bayes estimator [9] of $\mathrm{E}\{\theta_i|\boldsymbol{y}(1), y_i(2)\}$, which is derived under the assumption that $(\boldsymbol{y}(1), y_i(2))$ are independent. The naive Bayes estimator is

$$\hat{\theta}_i = \mu_\theta + \frac{y_i(1)\sqrt{\lambda_1} + y_i(2)\sqrt{\lambda(i,2)} - \mu_\theta(\lambda_1 + \lambda(i,2))}{\lambda_1 + \lambda(i,2) + \frac{\sigma^2}{\sigma_\theta^2}}. \tag{46}$$

In Fig. 1, we plot the MSE performance gain $g(\lambda)$, defined as

$$g(\lambda) = 10 \log \frac{\mathrm{var}(\theta_i|y_i(0))}{\mathrm{MSE}(\hat{\theta}_i)} \tag{47}$$

as a function of SNR, where Monte Carlo simulations were used to estimate the MSE of (46). We chose $Q = 4096$, $p \in \{1/10, 1/100\}$, and each point on the figure represents an average over θ based on 2000 realizations. Curves with crosses

and circles represent the optimal policy (ARAP) λ_A and the suboptimal policy λ_{so}, respectively. Note that ARAP yields better estimation performance compared to the suboptimal policy. The high gains in Fig. 1 reinforce the connection between our cost function (5) and estimation MSE.

B. Detection Postprocessing

1) Nonadaptive Detection: Consider the problem of correctly detecting whether cell i contains a target based on a sample $y_i(0)$. As before, we assume the samples follow the model (44) and that $\theta_i \sim \mathcal{N}\left(\mu_\theta, \sigma_\theta^2\right)$ are i.i.d. Thus, for an exhaustive search policy $y_i(0) \sim \mathcal{N}\left(\sqrt{\lambda_0}\mu_\theta I_i, \sigma_y^2\right)$, where $\sigma_y^2 = \sigma^2 + \lambda_0 \sigma_\theta^2$. Given $y_i(0)$, the measurement of pixel i, our goal is to decide between

$$H_0 : y_i(0) \sim \mathcal{N}(0, \sigma_y^2)$$
$$H_1 : y_i(0) \sim \mathcal{N}(\sqrt{\lambda_0}\mu_\theta, \sigma_y^2). \tag{48}$$

For a known σ_y^2, the uniformly most powerful test for this binary hypothesis testing problem is a likelihood ratio test (LRT). The performance of this nonadaptive LRT in terms of its receiver operating characteristic (ROC) curve is easily calculated. The power β of this level α LRT is

$$\beta = 1 - \Phi\left(\Phi^{-1}(1-\alpha) - \sqrt{\frac{\lambda_0 \mu_\theta^2}{\sigma^2 + \lambda_0 \sigma_\theta^2}}\right) \tag{49}$$

where $\Phi(\cdot)$ is the normal cumulative distribution function [10].

2) Adaptive Detection: Using Bayes rule, the likelihood function $f(\boldsymbol{y}(2), \boldsymbol{y}(1); \boldsymbol{I}_\Psi)$ equals

$$f(\boldsymbol{y}(2), \boldsymbol{y}(1)) = f(\boldsymbol{y}(2)|\boldsymbol{y}(1))f(\boldsymbol{y}(1)) \tag{50}$$

but, given $\boldsymbol{y}(1)$, the measurements at the second step are independent for different cells and thus

$$f(\boldsymbol{y}(2), \boldsymbol{y}(1)) = \prod_{i=1}^{Q} f(y_i(2)|\boldsymbol{y}(1))f(y_i(1)). \tag{51}$$

Therefore, the LRT statistic T_j is

$$T_j = \frac{\prod_{i=1}^{Q} f(y_i(2)|\boldsymbol{y}(1), I_i = 1)f(y_i(1)|I_i = 1)}{\prod_{i=1}^{Q} f(y_i(2)|\boldsymbol{y}(1), I_i = 0)f(y_i(1)|I_i = 0)} = $$
$$\frac{f(y_j(2)|\boldsymbol{y}(1), I_j = 1)f(y_j(1)|I_j = 1)}{f(y_j(2)|\boldsymbol{y}(1), I_j = 0)f(y_j(1)|I_j = 0)}. \tag{52}$$

For our model, we have $y_i(1) \sim \mathcal{N}\left(\sqrt{\lambda_1}\mu_\theta I_i, \sigma^2 + \lambda_1 \sigma_\theta^2\right)$ and, given $\boldsymbol{y}(1)$, the second step measurements are also Gaussian

$y_i(2) \sim \mathcal{N}\left(\sqrt{\lambda(i,2)}\mu_\theta I_i, \sigma^2 + \lambda(i,2)\sigma_\theta^2\right)$. Substituting these distributions into (52) provides the LRT shown in the equation at the bottom of the page, which simplifies to

$$T_i = \mu_\theta \left(\frac{y_i(1)\sqrt{\lambda_1}}{\sigma^2 + \lambda_1 \sigma_\theta^2} + \frac{y_i(2)\sqrt{\lambda(i,2)}}{\sigma^2 + \lambda(i,2)\sigma_\theta^2}\right)$$
$$- \frac{\mu_\theta^2}{2}\left(\frac{\lambda_1}{\sigma^2 + \lambda_1 \sigma_\theta^2} + \frac{\lambda(i,2)}{\sigma^2 + \lambda(i,2)\sigma_\theta^2}\right) \gtrless \gamma' \tag{53}$$

where $\gamma' = \log\gamma$. Note that for all cells where $\lambda(i,2) = 0$, this test is a function of $y_i(1)$ alone.

Next, we compare the theoretical ROC curve (49) to the empirical ROC curve calculated for the adaptive LRT (53) performed on the data pair $(\boldsymbol{y}(1), \boldsymbol{y}(2))$ using the optimal λ_A and the suboptimal λ_{so}. We conducted multiple runs for varying SNR levels and observed that with $\nu = 1$, (53) provides higher detection probability than the nonadaptive LRT having performance (49) for false alarm levels lower than 30%. At SNR values close to 0 dB, the difference between the two tests is negligible but increases with SNR. Note that, for very low false alarm levels, λ_A performs better than λ_{so} in terms of detection probability. However, for higher test levels, the suboptimal search policy λ_{so} yields better detection performances.

Results are presented in Fig. 2(a) and (b) for SNR = 10 dB, $Q = 1024$, and either $p = 0.1$ or $p = 0.01$. Monte Carlo simulations were conducted with $\nu = 1$, and each point on the figures represents an average over 2000 realizations. Detection probability was averaged over the entire ensemble and over all pixels inside and outside the ROI. At 10 dB, the ROC curves are very sharp and hence are plotted on a logarithmic scale. The solid curve represents [(49)] the nonadaptive LRT with equal energy allocation. Curves with crosses represent λ_A, while curves with circles represent λ_{so}. It is evident that the ROC curves of different tests have different slopes for low false-alarm values. Moreover, for high false-alarm values, no adaptive policy outperforms the exhaustive search policy. Fig. 2(b) zooms in to better illustrate the differences for $P_{fa} \in [0.005, 0.5]$. One can see that the optimal search policy has the best performances up to $\alpha = 3\%$, at which point the suboptimal policy yields higher detection probability. The exhaustive search policy outperforms both the adaptive methods for $\alpha > 30\%$.

Finally, we compare detection probability values, for a fixed false alarm rate, and estimation MSE gain $g(\cdot)$ in (47) as a function of ν. Results are shown in Fig. 3. The curve with triangle markers represents estimation MSE gain, and its corresponding values are indexed on the right axis of the figure. The other curves represent detection probability for a given test level, with the detection values indexed at the left axis. All curves are a

$$\frac{\exp\left\{-\frac{1}{2(\sigma^2 + \lambda(i,2)\sigma_\theta^2)}\left(y_i(2) - \mu_\theta\sqrt{\lambda(i,2)}\right)^2 - \frac{1}{2(\sigma^2 + \lambda_1 \sigma_\theta^2)}\left(y_i(1) - \mu_\theta\sqrt{\lambda_1}\right)^2\right\}}{\exp\left\{-\frac{y_i^2(1)}{2(\sigma^2 + \lambda_1 \sigma_\theta^2)} - \frac{y_i^2(2)}{2(\sigma^2 + \lambda(i,2)\sigma_\theta^2)}\right\}} \gtrless \gamma$$
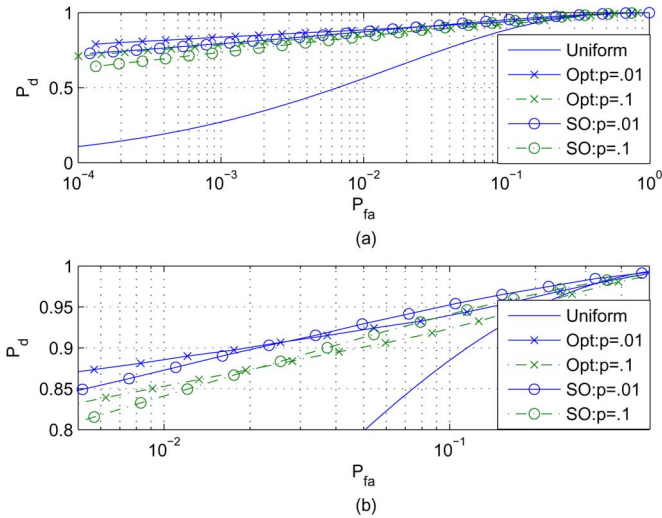
Fig. 2. ROC curves for the LRT tests based on an exhaustive search scheme and the two adaptive policy measurements scheme, for $p = 0.1$ and $p = 0.01$ and SNR of 10 dB. (a) shows the entire ROC curve while (b) zooms in on false alarm probability values less than 0.5. The simulation suggests that our adaptive search policies outperform an exhaustive search policy in terms of detection probability for any false alarm lower than 30%.
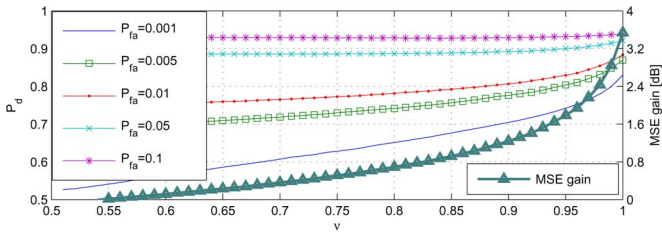

Fig. 3. Detection probability, for a fixed test level, and estimation MSE gain $g(\cdot)$ in (47) as a function of ν when SNR is 10 dB and $p = 0.01$. Note that the MSE gain values (curve with triangular markers) are given on the RHS of the figure. Since MSE gain is defined over the true ROI, it increases with ν. For this particular operating point, selecting $\nu = 1$ provides optimal performance for both detection and estimation tasks.
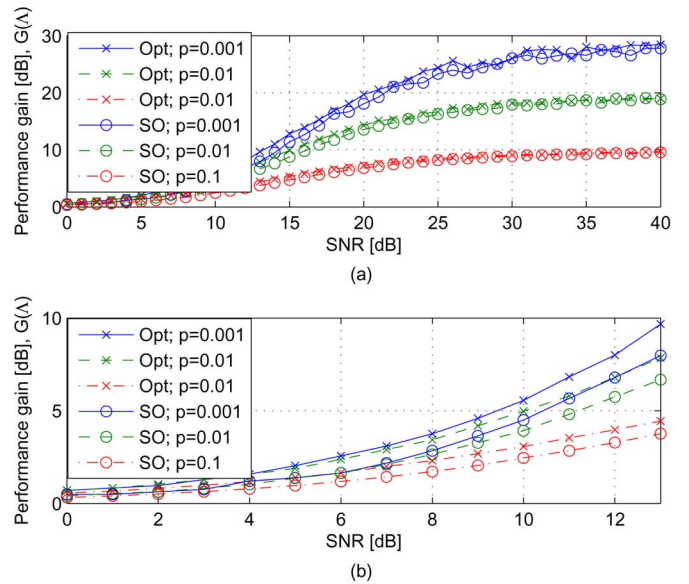

Fig. 4. The cost gain (15) compared to an exhaustive search for both our optimal and suboptimal energy allocation schemes. (a) shows that both algorithms converges to the asymptotic predicted gain, at $-10 \log p$. (b) enhances the difference between our two policies for SNR values in the range of 0–13 dB.
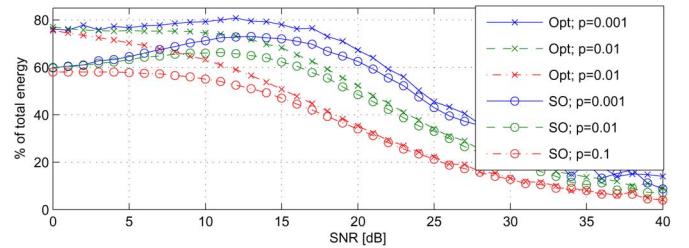

Fig. 5. The proportion of energy allocated at the first step for the two algorithms λ_A and λ_{so}. Curves correspond to prior probability values of 0.001, 0.01, and 0.1. As seen, the optimal search policy allocates more energy at the first step. However, for SNR > 25 dB, the two are essentially equivalent.

function of ν. For the selected operating point, it is clear that it is best to choose $\nu = 1$, since it maximizes both detection and estimation performance. In [7], we compare (53) to a test using a detection optimized measurement and show that λ_A with $\nu < 1$ is nearly optimal for detection.

C. Achievable Cost

As a final comparison between our two policies, we compare the average performance gain $G(\cdot)$, defined as the expected value of (15), achieved by the two search policies for $\nu = 1$. We chose $Q = 8192$ as the total number of cells and the sparsity values of $p \in \{1/10, 1/100, 1/1000\}$, i.e., a mean of roughly 800, 80, and 8 targets per realization, respectively. Results are shown in Fig. 4(a) and (b), where the curves with crosses and circles describe the expected gain $\mathrm{E}\{G(\Lambda_A)\}$ with optimal and with suboptimal $\mathrm{E}\{G(\Lambda_{so})\}$ allocation, respectively. Fig. 4(a) shows the behavior of the gain for the two policies for SNR values of 0–40 dB. Fig. 4(b) zooms in on SNR values of 0 to 13 dB. Each point on a graph represents 500 runs in a Monte Carlo simulation. As can be seen from Fig. 4(a), at extreme high or low SNR values, the performance gains of the two policies

coincide. Fig. 4(a) indicates that the gain converges to its theoretical limit given by iii). The largest gap in gain between the two algorithms is near the transition zone: SNR between 5–15 dB, and the gap is less than 2 [dB]. Evidently, the simpler suboptimal mapping rule does not significantly degrade performance gain.

Fig. 5 compares the percentage of effort $\lambda_1 Q / \lambda_T$ allocated in the first step for both policies. While for SNR values greater than 25 dB the curves overlap, this is not the case for low SNR values. As measurement quality decreases, ARAP invests more energy at the first step. Considering the difference between the two policies, this result makes sense: after the first step, ARAP ignores all cells with posterior probability values lower than some threshold. Decision errors at that stage can no longer be compensated by increasing energy allocation at the next stage. Hence, more effort has to be allocated to the first step to reduce decision errors, i.e., improve the agreement between $\widehat{\Psi}$ and Ψ. On the other hand, the suboptimal mapping invests energy in all Q cells at the second step. Thus, it has a chance to compensate for poor estimated posterior probability values. As a general remark, we note that the incentive to use our methods increases

with SNR (e.g., Fig. 1). Therefore, differences between the two policies in the low SNR regime are of lesser importance.

V. APPLICATION—DETECTING AND ESTIMATING AN ROI IN A RADAR IMAGING SYSTEM

Consider the task of imaging a large area using radar. Such a task is routinely performed by air traffic control (ATC) radar, early warning radar (EWR), and surface movement radar (SMR), where the different systems are looking for small targets (aircrafts, missiles, people, etc.) over a large domain. A radar beam is either steered mechanically (ATC, SMR) or electronically (EWR) using a predefined grid pattern designed to detect a minimal target profile. The target profile dictates scan parameters such as beam pattern, grid spacing, and dwell time or the number of pulses transmitted at each grid point. Such applications are particularly suitable for ARAP since no a priori knowledge is required and resulting images are sparse with respect to targets. The following is motivated by wide-band SMR imaging system for ground movement monitoring in airports [11].

The fixed budget at hand is the overall scan time. Assume that an image is to be formed, i.e., scan the entire domain, every λ_T seconds. The scan time $\lambda_T = Q\lambda_0$, where Q is the total number of grid points and λ_0 is dwell time at each point, determined by the minimal target profile or minimal SNR. We assume that measurement quality improves as the dwell time increases. Therefore, we address the question of how to best utilize the resource available to us, i.e., where to dwell and for how long to dwell the radar beam. Problems concerning where to point a sensor are discussed in [12] and [13]. However, by allowing an additional continuous degree of freedom for the varying dwell time, we consider a different set of feasible policies. We propose the following two-stage search policy based on ARAP: at the first stage, perform an exhaustive search with equal dwell time $\lambda_1 = \alpha\lambda_T/Q$ over all cells yielding a preliminary image $\boldsymbol{y}(1)$ and an allocation set $\widehat{\Psi}$. In the second stage, let ARAP allocate the remaining $(1-\alpha)\lambda_T$ seconds in a nonuniform manner to the different cells in $\widehat{\Psi}$ via (37), yielding a nonuniform number of pulses transmitted at each cell. Finally, we combine both measurements to form an image on which we detect the ROI and estimate its content. The competing strategy performs an exhaustive search with equal dwell time $\lambda_0 = \lambda_T/Q$ at each cell to form the estimated image. Our goal was not to simulate a w-band imaging system but rather to show the applicability of ARAP to such systems. Therefore, we used a SAR image, taken from the Sandia National Laboratories website, as an example of a "sparse" image. The image displays two columns of tanks in a field, and its sparsity ratio is $p < 0.01$. We simulated several noise models to evaluate the performance of ARAP. The tested noise models included: Speckle, Rayleigh, Swerling II, and Gaussian. Here we only present simulation results using the Swerling model II (see [14, pp. 57]) as they are representative of all other models tested. Note that the Swerling II is a non-Gaussian model that is mismatched to the model assumed in the ARAP derivation, and therefore it represents a good test of robustness to model mismatch.

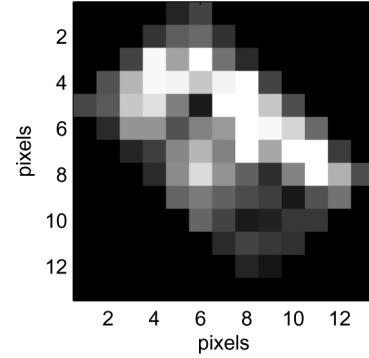[1]http://www.sandia.gov/RADAR/images/rtv_tanks_9in.jpg.

Fig. 6. The above (13×13) tank template was used as a matched filter to filter the noisy data X_1 and generate $\boldsymbol{y}(1)$.

Let X denote the original image and let $\boldsymbol{x}$ be a lexicographic ordering of X. The image after the first stage (λ_1 pulses at all pixels) is modeled as

$$\boldsymbol{x}_1 = \frac{1}{\lambda_1} \sum_{j=1}^{\lambda_1} \boldsymbol{z}(j) \qquad (54)$$

where $z_i(j)$ is an exponential random variable, i.e., $f_{Z_i}(z) = 1/\boldsymbol{x}(i) \exp\{-z/\boldsymbol{x}(i)\}$, denoted $z_i(j) \sim \mathrm{Exp}(\boldsymbol{x}(i))$, and $\boldsymbol{x}(i)$ is the ith element in the lexicographic ordering of the original image X. A tank template shown in Fig. 6 was applied as a matched filter to the noisy image X_1 yielding $\widetilde{X}_1$. The input to the ARAP algorithm $\boldsymbol{y}(1) = (\widetilde{\boldsymbol{x}}_1 - \overline{x}_1)/\sqrt{(\widetilde{\boldsymbol{x}}_1 - \overline{x}_1)'(\widetilde{\boldsymbol{x}}_1 - \overline{x}_1)/Q}$ was the variance normalized version of $\widetilde{X}_1$, where $\overline{x}_1 = 1/Q \sum_{i=1}^{Q} \widetilde{\boldsymbol{x}}_1(i)$. ARAP was used to generate a search policy for the second step $\lambda_A(i,2)$ via (37). All indexes i with $\lambda_A(i,2) < \lambda_1$ were set to zero, and their cumulative search effort was redistributed among the rest of the cells in a proportional manner. The radar return $\boldsymbol{x}_2$ was modeled as

$$\boldsymbol{x}_2(i) = \frac{1}{\lambda_A(i,2)} \sum_{j=1}^{\lambda_A(i,2)} z_i(j), \quad i \in \widehat{\Psi}. \qquad (55)$$

This resulted in a nonuniform variance for different cells in $\widehat{\Psi}$ with $\mathrm{var}(z_2(i)) = \boldsymbol{x}^2(i)/\lambda_A(i,2)$. Note that $\boldsymbol{x}_2 = 0$ for all $i \in \widehat{\Psi}^c$, i.e., $\boldsymbol{I}_{\widehat{\Psi}}$ plays the role of the energy allocation indicator.

We considered several suboptimal linear estimators of $\boldsymbol{x}$ based on $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ that performed comparably. The estimator presented here $\widehat{X}$ was defined via $\widehat{\boldsymbol{x}}$ as

$$\widehat{\boldsymbol{x}} = \left(\frac{\boldsymbol{x}_1 \odot I\{i \in \widehat{\Psi}^c\}}{\sqrt{1 + 1/\lambda_1}} + \boldsymbol{x}_2 \odot \mathrm{vec}\left\{ \frac{1}{\sqrt{1 + 1/\lambda(i,2)}} \right\} \right) \odot \boldsymbol{s} \qquad (56)$$

where

$$\mathrm{vec}\{r_i\} = [r_1, r_2, \ldots, r_Q]',$$
$$\boldsymbol{s}(i) = \left(\sqrt{1 + 1/\lambda_1} + \sqrt{1 + I\{i \in \widehat{\Psi}\}/\lambda(i,2)} \right)^{-1},$$

and $\odot$ represent an element by element multiplication. The estimator (56) is compared to an image reconstructed from a single
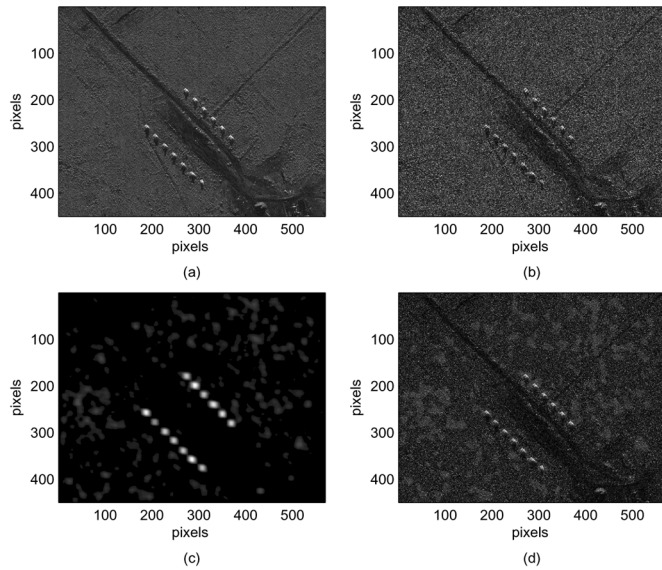
Fig. 7. Radar imaging example, $\lambda_0 = 3$. (a) Original image. (b) Image reconstructed using an exhaustive search. (c) Effort allocation using ARAP at the second stage ($\lambda_1 = 2$). (d) Image resulted from (56) using ARAP.
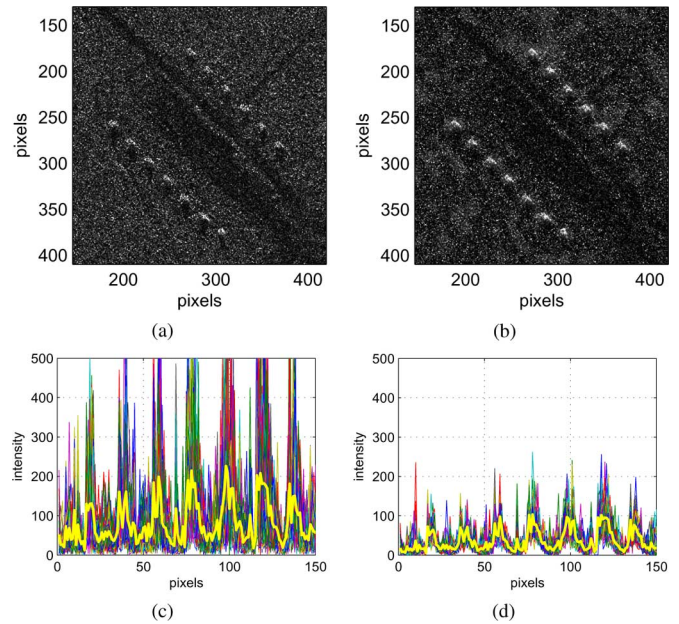

Fig. 8. Radar imaging example, $\lambda_0 = 2$. (a) Image reconstructed using an exhaustive search, targets are not easily identifiable. (b) Image resulted from (56) using ARAP ($\lambda_1 = 1$). (c) and (d) compare a 1-D profile going through the targets on the lower left column for 100 different realizations. (c) Profiles of images reconstructed from an exhaustive search. (d) Profiles of images reconstructed using ARAP. The bold bright line on both figures represent the mean profile of the different realizations. Evidently, variations of profiles of images due to ARAP are much smaller compared to variations of profiles of images resulted from an exhaustive scan, with variance reduction exceeding 10 dB.

exhaustive scan $\boldsymbol{x}_U$ using λ_0 pulses at each cell via (54) replacing λ_1 with λ_0.

Results are presented in Figs. 7 and 8 for λ_0 values of three and two pulses, respectively. Fig. 7(a) shows the original image, and Fig. 7(b) and (d) shows a single realization of images reconstructed using exhaustive search and ARAP via (56), respectively. Fig. 7(c) shows the effort allocated by ARAP at the second stage for that specific realization. The bright area over the true ROI corresponds to more than 40 pulses per pixel. Although all targets are identifiable in Fig. 7(b), they appear more clearly in Fig. 7(d). Fig. 8 focuses on the ROI to demonstrate the superiority of ARAP compared to the exhaustive search policy. Fig. 8(a) and (b) shows a single realization of the two search methods, exhaustive and ARAP, for $\lambda_0 = 2$ ($\lambda_1 = 1$ for ARAP), respectively, while Fig. 8(c) and (d) displays a one-dimensional (1-D) profile, going through the left columns of tanks, of 100 different realizations of each policy, respectively. Clear difference in variations of image profiles reconstructed using ARAP compared to the images reconstructed from an exhaustive search is evident. This difference corresponds to more than 10 dB reduction in variance in the ROI.

This illustrative example demonstrates the potential utility of our method in radar imaging. Note that energy allocation is equivalent to dwell-time allocation in this example.

VI. CONCLUSIONS AND FUTURE WORK

We introduced a novel convex cost function and showed that its minimization relates to minimizing error probability or estimation MSE over an unknown ROI. A closed-form solution for the second stage in a two-stage optimal search policy was provided, and numeric search for the first step was presented. A closed-form approximation to the two-step minimization problem was shown to perform comparably to the optimal solution. In the limit of high SNR, the performance of the optimal and approximated algorithms converges to the ideal

omniscient limit. For the detection task, the two search policies introduced outperformed the one-step exhaustive measurement scheme when the false alarm is less than 30%. For estimation, comparing the MSE of estimated values within the ROI, our adaptive search policies dominate the exhaustive search policy. The search policy is parameterized by ν, which varies from 1/2 to 1 and controls the energy allocated within the ROI. An offline lookup table can be generated for the optimal ν in terms of the sparseness p of the target in the data and SNR. Finally, an illustrative example of our method for radar imaging was presented.

This approach is applicable to tumor detection where a cluster of calcification may appear around the lesion. In this case, multiscale hypothesis-testing methods presented in [15] may be relevant. Frakt *et al.* deal with anomaly detection once measurements, at a fine resolution, have been acquired. A goal would be to generate fine-resolution measurements only where they are most informative. Another interesting area of application is to compressive sensing. Works such as [16] and [17] consider the problem of sampling a sparse medium via an arbitrary affine transformation. In cases of sparse signals, complete reconstruction of the underlying signal can be accomplished with only a few samples. This exploitation of sparsity might benefit from the Bayesian methods proposed in this paper.

APPENDIX

Let $\boldsymbol{Y}(t) \in \mathcal{Y}$ be a $(Q \times 1)$ RV with a probability density function (pdf) $p_{\boldsymbol{Y}(t)}(\boldsymbol{y}) > 0$, for all $\boldsymbol{y} \in \mathcal{Y}$ and $t \in \{1, 2, \dots, T\}$, representing random observations. Let

$\boldsymbol{I}_\Psi = [I_1, I_2, \ldots, I_Q]'$ be an RV with $\Pr(I_i = 1) = p_i$. Let $Y_1^t = \{\boldsymbol{Y}(1), \boldsymbol{Y}(2), \ldots, \boldsymbol{Y}(t)\}$ be a collection of all observations up to time t. Define $x(i,t) : Y_1^{t-1} \to \mathbb{R}_+$; then, for some $\nu \in (1/2, 1]$, our goal is to find $\{\hat{x}(i,t)\}$ with $i = 1, \ldots, Q$ and $t = 1, \ldots, T$, such that

$$\hat{x}(i,t) = \arg\min_{x(i,t)} \mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{\nu I_i + (1-\nu)(1-I_i)}{\sum_{t=1}^{T} x(i,t)} \right\}$$

where expectation is taken with respect to $\boldsymbol{I}_\Psi$ and Y_1^{T-1}, subject to $\sum_{i=1}^{Q} \sum_{t=1}^{T} x(i,t) = X$.

A. The Case of $T = 1$

For $T = 1$, our cost function has the following form:

$$\hat{x}(i) = \arg\min_{x(i)} \mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{\nu I_i + (1-\nu)(1-I_i)}{x(i)} \right\}$$

subject to $\sum_{i=1}^{Q} x(i) = X$, with the expectation taken with respect to $\boldsymbol{I}_\Psi$. Note that $\mathrm{E}\{I_i\} = p_i$, so $\hat{x}_i$ can be derived using Lagrange multiplier, i.e., finding the minimizer of

$$L(\boldsymbol{x}, \lambda) = \sum_{i=1}^{Q} \frac{\nu p_i + (1-\nu)(1-p_i)}{x(i)} + \lambda \left(\sum_{i=1}^{Q} x(i) - X \right).$$

Taking derivatives and setting them equal to zero yields

$$\hat{x}(i) = \frac{X \sqrt{\nu p_i + (1-\nu)(1-p_i)}}{\sum_{j=1}^{Q} \sqrt{\nu p_j + (1-\nu)(1-p_j)}}.$$

B. The Case of $T = 2$

Consider the following problem, let $x(i,1) = x_1 > 0$ for all i; then $\sum_{i=1}^{Q} x(i,1) = Q x_1$. Our goal is minimize the cost, i.e., find

$$\hat{x}(i,2) = \arg\min_{x(i,2)} \mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{\nu I_i + (1-\nu)(1-I_i)}{x_1 + x(i,2)} \right\} \quad (57)$$

subject to

$$\sum_{i=1}^{Q} x(i,2) = X - Q x_1 \geq 0, \quad (58)$$

and $x(i,2) : \boldsymbol{Y}(1) \to \mathbb{R}_+$. For brevity, let $x_i(\boldsymbol{Y}) = x(i,2)$, and note that in (57) expectation is taken with respect to $\boldsymbol{I}_\Psi$ and $\boldsymbol{Y}(1)$. Using iterated expectation, we obtain

$$\mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{\nu I_i + (1-\nu)(1-I_i)}{x_1 + x(i,2)} \right\}$$
$$= \mathrm{E} \left\{ \sum_{i=1}^{Q} \mathrm{E} \left\{ \frac{\nu I_i + (1-\nu)(1-I_i)}{x_1 + x_i(\boldsymbol{Y})} \Big| \boldsymbol{Y}(1) \right\} \right\}. \quad (59)$$

Given $\boldsymbol{Y}(1)$, the denominator is deterministic and expectation can be applied to the numerator; therefore (59) becomes

$$\mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{\nu \Pr(I_i = 1 | \boldsymbol{Y}(1)) + (1-\nu)(1 - \Pr(I_i = 1 | \boldsymbol{Y}(1)))}{x_1 + x_i(\boldsymbol{Y})} \right\}.$$

Defining $p_{I_i | \boldsymbol{Y}} = \Pr(I_i = 1 | \boldsymbol{Y}(1))$ and $W_i = \nu p_{I_i | \boldsymbol{Y}} + (1-\nu)(1 - p_{I_i | \boldsymbol{Y}})$, we use a Lagrange multiplier to minimize (60) as shown at the bottom of the page, where $w_i = W_i(\boldsymbol{y})$ is a realization of the random variable W_i. Since $f_{\boldsymbol{Y}}(\boldsymbol{y})$ is strictly positive, define $\lambda(\boldsymbol{y}) = \lambda'(\boldsymbol{y})/f_{\boldsymbol{Y}}(\boldsymbol{y})$; then, $\hat{x}_i$, the minimizer of (60) is also the minimizer of

$$\sum_{i=1}^{Q} \int_{\mathcal{Y}} \left[\frac{w_i}{x_1 + x_i(\boldsymbol{y})} + \lambda(\boldsymbol{y}) x_i(\boldsymbol{y}) \right] d\boldsymbol{y}.$$

Note that our problem has translated to minimizing a separable sum of integrals of positive integrands. Hence, finding

$$\hat{x}_i(\boldsymbol{Y}) = \arg\min_{x_i(\boldsymbol{Y})} \left\{ \frac{w_i}{x_1 + x_i(\boldsymbol{Y})} + \lambda(\boldsymbol{Y}) x_i(\boldsymbol{Y}) \right\}$$

suffice. Solving for $x_i(\boldsymbol{y})$ given the Lagrange multiplier $\lambda(\boldsymbol{Y})$ yields

$$x_i(\boldsymbol{Y}) = \begin{cases} \sqrt{\frac{w_i}{\lambda(\boldsymbol{Y})}} - x_1, & \sqrt{\frac{w_i}{\lambda(\boldsymbol{Y})}} > x_1 \\ 0, & \text{otherwise} \end{cases} \quad (61)$$

$$\mathrm{E} \left\{ \sum_{i=1}^{Q} \frac{W_i}{x_1 + x_i(\boldsymbol{Y})} \right\}$$
$$+ \int_{\mathcal{Y}} \lambda'(\boldsymbol{y}) \left(\sum_{i=1}^{Q} x_i(\boldsymbol{y}) - (X - Q x_1) \right) d\boldsymbol{y}$$
$$= \sum_{i=1}^{Q} \int_{\mathcal{Y}} \left[\frac{w_i}{x_1 + x_i(\boldsymbol{y})} f_{\boldsymbol{Y}}(\boldsymbol{y}) + \lambda'(\boldsymbol{y})(x_i(\boldsymbol{y}) - (X - Q x_1)) \right] d\boldsymbol{y} \quad (60)$$

which can be also written as $x_i(\boldsymbol{Y}) = (\sqrt{w_i/\lambda(\boldsymbol{Y})} - x_1)I(\lambda(\boldsymbol{Y}) < w_i/x_1^2)$, where $I(\cdot)$ is an indicator function. Utilizing the constraint $\sum_i x_i(\boldsymbol{Y}) = X - Qc_1$, we obtain

$$\frac{1}{\sqrt{\lambda(\boldsymbol{Y})}} \sum_{i=1}^{Q} \sqrt{w_i} I\left(\lambda(\boldsymbol{Y}) < \frac{w_i}{x_1^2}\right)$$
$$= X - x_1 \left(Q - \sum_{i=1}^{Q} I\left(\lambda(\boldsymbol{Y}) < \frac{w_i}{x_1^2}\right)\right). \quad (62)$$

Note that

$$Q = \sum_{i=1}^{Q} I\left(\lambda(\boldsymbol{Y}) < \frac{w_i}{x_1^2}\right) + \sum_{i=1}^{Q} I\left(\lambda(\boldsymbol{Y}) \geq \frac{w_i}{x_1^2}\right) \quad (63)$$

substituting (63) into (62) and rearranging yields

$$\sqrt{\lambda(\boldsymbol{Y})} = \frac{\sum_{i=1}^{Q} \sqrt{w_i} I\left(\lambda(\boldsymbol{Y}) < \frac{w_i}{x_1^2}\right)}{X - x_1 \sum_{i=1}^{Q} I\left(\lambda(\boldsymbol{Y}) \geq \frac{w_i}{x_1^2}\right)}. \quad (64)$$

Next, use $\tau : \mathcal{X} \to X$ defined in (31) such that $w_{\tau(1)} \leq w_{\tau(2)} \leq \cdots \leq w_{\tau(Q)}$.

1) Case $\lambda(Y)x_1^2 < \omega_{\tau(1)}$: If $\lambda(\boldsymbol{Y})x_1^2 < w_{\tau(1)}$, for all $\boldsymbol{Y}$, then $\lambda(\boldsymbol{Y})x_1^2 < w_{\tau(i)}$ for all i and $\sqrt{\lambda(\boldsymbol{Y})} = \sum_{i=1}^{Q} \sqrt{w_{\tau(i)}}/X$. For which case, the cost minimizer is achieved at

$$x_{\tau(i)}(\boldsymbol{Y}) = \frac{X\sqrt{w_{\tau(i)}}}{\sum_{j=1}^{Q} \sqrt{w_j}} - x_1$$

and, for all i

$$\frac{X}{x_1} \geq \frac{\sum_{j=1}^{Q} \sqrt{w_j}}{\sqrt{w_{\tau(i)}}}. \quad (65)$$

2) Case $\lambda(Y)x_1^2 \geq \omega_{\tau(Q)}$: For the other extreme, we have $\lambda(\boldsymbol{Y})x_1^2 \geq w_{\tau(i)}$ for all i, which in turn forces $\lambda(\boldsymbol{Y}) = 0$, due to (64). This can only hold if both $\nu = 1$ and $p_{I_i|\boldsymbol{Y}} = w_{\tau(i)} = 0$, for all i. If this is the case, then for any choice of $x_i(\boldsymbol{Y})$, the overall cost is zero.

3) Case $\omega_{\tau(k)} \leq \lambda(Y)x_1^2 < \omega_{\tau(k+1)}$, for Some $k \in \{1, 2, \ldots, q-1\}$: The interesting case is when $\lambda(\boldsymbol{Y})$ has some intermediate value, i.e.,

$$\frac{\sqrt{w_{\tau(k)}}}{x_1} \leq \frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{X - kx_1} < \frac{\sqrt{w_{\tau(k+1)}}}{x_1}. \quad (66)$$

Since all the terms in (66) are positive, we rewrite the inequality as

$$\frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k+1)}}} + k < \frac{X}{x_1} \leq \frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k)}}} + k. \quad (67)$$

Define

$$\text{LB}(k) = \frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k+1)}}} + k$$

$$\text{UB}(k) = \frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k)}}} + k. \quad (68)$$

To show that (67) makes sense, we need to show that for any X of interest, there exists a unique $k = k_0 \in \{1, 2, \ldots, Q-1\}$ such that $X/x_1 \in (\text{LB}(k_0), \text{UB}(k_0)]$. To do so, we show that following: i) $\text{LB}(k) \leq \text{UB}(k)$, ii) $\text{LB}(k-1) = \text{UB}(k)$, and iii) $X/x_1 \in (\text{LB}(Q-1), \text{UB}(1)]$. Start with i)

$$\frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k+1)}}} + k - \left(\frac{\sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}}}{\sqrt{w_{\tau(k)}}} + k\right)$$
$$= \left(\frac{\sqrt{w_{\tau(k)}} - \sqrt{w_{\tau(k+1)}}}{\sqrt{w_{\tau(k)}w_{\tau(k+1)}}}\right) \sum_{i=k+1}^{Q} \sqrt{w_{\tau(i)}} \leq 0 \quad (69)$$

since $w_{\tau(k)} \leq w_{\tau(k+1)}$. Note that $w_{\tau(k)} = w_{\tau(k+1)}$ implies that the interval $(\text{LB}(k), \text{UB}(k)]$ is empty. To prove ii), we substitute $(k-1)$ for k in the expression for $\text{LB}(k)$ in (68). This yields

$$\text{LB}(k-1) - \text{UB}(k) = \frac{\sqrt{w_{\tau(k)}}}{\sqrt{w_{\tau(k)}}} - 1 = 0. \quad (70)$$

Combining i) and ii) provides $\text{LB}(k) \leq \text{UB}(k) = \text{LB}(k-1)$, i.e., $\text{LB}(k)$ is decreasing in k. Finally, we need to show $X/x_1 \in (\text{LB}(Q-1), \text{UB}(1)]$. First, note from (65) that $X/x_1 > \text{UB}(1)$ belongs to *Case 1,* for which a separate solution was provided. Secondly, from (68), we obtain $\text{LB}(Q-1) = Q$. If $Qx_1 = X$, then $x_i(\boldsymbol{Y}) = 0$, for all i, due to the constraint (58). Thus, for a two-step problem, we must have $X/x_1 > Q = \text{LB}(Q-1)$. Hence, $X/x_1 \in (\text{LB}(Q-1), \text{UB}(1)]$ as required. Together i)–iii) prove the existence and uniqueness of the solution, since

1. $\dfrac{X}{x_1} \in (\text{LB}(Q-1), \text{UB}(1)]$

2. $\bigcup_{k=1}^{Q} (\text{LB}(k), \text{UB}(k)] = (\text{LB}(Q-1), \text{UB}(1)]$

3. $(\text{LB}(i), \text{UB}(i)] \cap (\text{LB}(j), \text{UB}(j)] = \emptyset \ \forall \ i \neq j$.

To conclude this section, we note that $\hat{x}_i(\boldsymbol{Y})$, the minimizer of (57), is given by (34), replacing x_1 for λ_1 and $\hat{x}_{\tau(i)}(\boldsymbol{Y})$ for $\lambda(\tau(i), 2)$. A general solution for the case of a nonuniform prior distribution, which is also a water filling algorithm, is given in [7].

REFERENCES

[1] E. Posner, "Optimal search procedures," *IEEE Trans. Inf. Theory*, vol. IT-9, pp. 157–160, Jul. 1963.
[2] I. M. Johnstone and B. W. Silverman, "Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences," *Ann. Statist.*, vol. 32, pp. 1594–1649, 2004.

[3] D. P. Wipf and B. D. Rao, "Sparse Bayesian learning for basis selection," *IEEE Trans. Signal Process.*, vol. 52, no. 8, pp. 2153–2164, Aug. 2004.

[4] R. Castro, R. Willett, and R. Nowak, "Faster rates in regression via active learning," presented at the Neural Inf. Process. Syst. Conf. (NIPS) 2005, Vancouver, BC, Canada, Dec. 2005.

[5] R. Rangarajan, R. Raich, and A. O. Hero, III, "Optimal sequential energy allocation for inverse problems," *IEEE J. Sel. Topics Signal Process.*, vol. 1, pp. 67–78, Jun. 2007.

[6] A. O. Hero, III, D. A. Castañón, D. Cochran, and K. Kastella, *Foundations and Applications of Sensor Management*. Boston, MA: Springer, 2007.

[7] E. Bashan, R. Raich, and A. O. Hero, III, "Adaptive sampling: Efficient search schemes under resource constraints," Communications and Signal Processing Lab, Univ. of Michigan, Ann Arbor, MI, 2007.

[8] S. Boyd and L. C. O. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[9] B. Wray, "Learning classification rules using Bayes," in *Proc. 6th Int. Workshop Machine Learn.*, San Francisco, CA, 1989, pp. 94–98.

[10] S. M. Kay, *Fundamentals of Statistical Signal Processing*. Upper Saddle River, NJ: Prentice-Hall, 1993, vol. 1, Estimation Theory.

[11] F. Mauro, G. Galati, and M. Naldi, "A novel W-band radar for airport traffic monitoring: implementation, first operational results, perspectives," in *Proc. 4th Int. Kharkov Symp. Phys. Eng. Millimeter Sub-Millimeter Waves 2001*, 2001, vol. 1, pp. 27–32.

[12] K. Kastella, "Discrimination gain to optimize detection and classification," *IEEE Trans. Syst., Man, Cybern. A*, vol. 27, pp. 112–116, Jan. 1997.

[13] C. Kreucher, K. Kastella, and A. O. Hero, III, "Sensor management using an active sensing approach," *Signal Process.*, vol. 85, no. 3, pp. 607–624, Mar. 2005.

[14] G. P. Kulemin, *Millimeter-Wave Radar Targets and Clutter*. Norwood, MA: Artech House, 2003.

[15] A. B. Frakt, W. C. Karl, and A. S. Willsky, "A multiscale hypothesis testing approach to anomaly detection and localization from noisy tomographic data," *IEEE Trans. Image Process.*, vol. 7, pp. 825–837, Jun. 1998.

[16] M. B. Wakin, J. N. Laska, M. F. Duarte, D. Baron, S. Sarvotham, D. Takhar, K. F. Kelly, and R. G. Baraniuk, "An architecture for compressive imaging," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2006, pp. 1273–1276.

[17] M. B. Wakin and R. G. Baraniuk, "Random projections of signal manifolds," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, May 2006, vol. 5.

Eran Bashan was born in Tel Aviv, Israel, in 1971. He received the B.Sc. degree in electrical and computer engineering from Ben-Gurion University, Beer-Sheva, Israel, in 2000, the M.Sc. degree from Tel-Aviv University, Tel Aviv, Israel, in 2004, and the Ph.D. degree from the University of Michigan, Ann Arbor, in 2008 both in electrical engineering (systems).

He is currently a Postdoctoral Fellow at the University of Michigan. His research focus is on adaptive and efficient resource-allocation policies with application for search. Other interests lie in detection, estimation theory, and sequential design, with applications to medical imaging modalities and target detection.

Raviv Raich (S'98–M'04) received the B.Sc. and M.Sc. degrees from Tel Aviv University, Tel-Aviv, Israel, in 1994 and 1998, respectively, and the Ph.D. degree from Georgia Institute of Technology, Atlanta, in 2004, all in electrical engineering.

Between 1999 and 2000, he was a Researcher with the communications Team, Industrial Research, Ltd., Wellington, New Zealand. From 2004 to 2007, he was a Postdoctoral Fellow with the University of Michigan, Ann Arbor. Since fall 2007, he has been an Assistant Professor in the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis. His main research interest is in statistical signal processing with specific focus on manifold learning, sparse signal reconstruction, and adaptive sensing. His other research interests lie in the area of statistical signal processing for communications, estimation and detection theory, and machine learning.
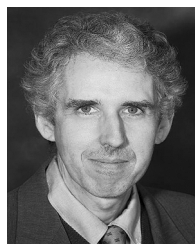
Alfred O. Hero (S'790–M'84–SM'96–F'98) received the B.S. degree (*summa cum laude*) from Boston University, Boston, MA, in 1980 and the Ph.D. degree from Princeton University, Princeton, NJ, in 1984, both in electrical engineering.

Since 1984, he has been with the University of Michigan, Ann Arbor, where he is a Professor in the Department of Electrical Engineering and Computer Science and, by courtesy, in the Department of Biomedical Engineering and the Department of Statistics. He has held visiting positions at the Massachusetts Institute of Technology (2006); Boston University; I3S University of Nice, Sophia-Antipolis, France (2001); Ecole Normale Supérieure de Lyon (1999); Ecole Nationale Supérieure des Télécommunications, Paris, France (1999); Scientific Research Labs of the Ford Motor Company, Dearborn, MI (1993); Ecole Nationale Supérieure des Techniques Avancees (ENSTA); Ecole Supérieure d'Électricité, Paris, France (1990), and M.I.T. Lincoln Laboratory (1987–1989). His recent research interests include the areas of inference for sensor networks, adaptive sensing, bioinformatics, inverse problems, and statistical signal and image processing.

Dr. Hero is a member of Tau Beta Pi, the American Statistical Association (ASA) , the Society for Industrial and Applied Mathematics (SIAM), and the U.S. National Commission (Commission C) of the International Union of Radio Science (URSI). He has received an IEEE Signal Processing Society Meritorious Service Award (1998), an IEEE Signal Processing Society Best Paper Award (1998), an IEEE Third Millennium Medal, and a 2002 IEEE Signal Processing Society Distinguished Lectureship. He was President of the IEEE Signal Processing Society (2006–2007) and during his term served on the TAB Periodicals Committee (2006). He is currently a member of the TAB Society Review Committee (2008–2010).