18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Adaptive Search for Multi-Class Targets with Heterogeneous Importance

Beipeng Mu, Aero. Astro., MIT, Cambridge, MA 02139, mubp@mit.edu
Gregory E. Newstadt, Google Pittsburgh, Pittsburgh, PA 15206, gregoryn@google.com
Dennis L. Wei, IBM Research, Yorktown Heights, NY 10598, dwei@us.ibm.com
Alfred O. Hero, EECS, U. of Michigan, Ann Arbor,MI 48109, hero@umich.edu
Jonathan P. How, Aero. Astro., MIT, Cambridge, MA 02139, jhow@mit.edu

***Abstract*—In sparse target detection problems, it has been shown that significant gains can be achieved by adaptive sensing. We generalize previous work on adaptive sensing to (a) include targets of multiple classes with different levels of mission importance and (b) account for multiple sensor models. New optimization policies are developed to simultaneously locate, classify and estimate a sparse number of targets with limited resource budget. More specifically, three sensor models are considered: global adaptive (GA) sensor that allocates different amounts of resource to each location in the space; global uniform (GU) sensor that allocates resources uniformly across the scene; and local adaptive (LA) sensor that allocates fixed amount of resources to a subset of locations. Based on the sensor model, we propose 3 policies: GA policy that uses a GA sensor; LA policy that uses only LA sensor; and GU/LA policy that uses a mixture of GU and LA sensors. The performances of proposed policies with these sensor models are compared numerically with a baseline policy that allocates resources uniformly and an oracle policy with known target locations. Results indicate that the GA policy performs closely to the oracle policy with sufficient resources, and the GU/LA policy performs similarly to that the GA policy but it is cheaper and more easily implementable.**

I. Introduction

This work considers localization, classification, and estimation of targets from observations taken sequentially and adaptively. In particular, we focus on the regime where (a) the number of targets is sparse compared to the size of the scene, and (b) some of the targets are more important than others. For example, in search-and-rescue missions, detection of survivors has significantly higher mission importance than detection of other features in the environment. Similarly, a radar operator may be more interested in detecting/tracking a tank rather than a car, though both might sparsely populate the scene.

Viewing the targets as sparse signals, adaptive localization and estimation use past observations to shape future measurements of the scene [1]–[3], which can result in stronger signal-to-noise ratios. The performance gain occurs by adaptively focusing the majority of sensing resources only on the dimensions of the signal that contain targets. Applications where adaptive sensing has been used include image acquisition/compression, spectrum sensing, agile radars, and medical imaging [2]–[7].

Previous work has shown that adaptive sensing problems can be formulated as partially observable Markov decision processes (POMDPs) [8]–[11]. Applications include surveillance [8,9] and robot coordination [11]. However, the complexity of POMDP solutions grows exponentially with the number of targets and sensors, making them generally intractable for the size of problems that we consider here. This scaling is further complicated by coupling of sensors, combination of continuous signals and actions, and heterogeneity of sensors/targets. In this work, we provide approximate solutions that are both tractable and also perform very well in comparison to baseline policies (lower bound) and oracle policies (upper bound).

Previous work in adaptive sensing [1]–[3] considered only the two-class problem, where a target is either absent or present. In many applications, such as surveillance and search-and-rescue missions, targets may have different classes with varying importance to the mission. In this setting, detection-based methods may waste resources on unimportant targets and thus suffer performance losses for the important ones. This work explicitly accounts for multi-class targets with different mission importance.

Previous work also depends on the availability of an agile sensor that can allocate sensing resources to any combination of locations in the scene at the same time and with potentially different effort (i.e. dwell time or energy) [2,6]. In some cases we may only have access to sensors with restricted agility. This work explicitly considered more realistic sensor models: a global uniform sensor that allocates the same amount everywhere, or a local adaptive sensor that allocates sensing resources only to a limited set of locations. In this paper, we analyze the performance differences between these types of sensors.

There is relatively limited work for planning with multiple types of sensors and tasking agents, particularly when the size of the scene is large. Previous work [12,13] proposed a planning algorithm for a team of heterogeneous sensors with the goal of maximizing mission importance. This algorithm showed that tasking agents achieved significantly better performance gains when target exploration included mission importance in the planning stage, as compared to an exploration policy that only depended on target uncertainty. But it was assumed that the number and location of targets were known a priori. In this work, we relax this assumption and simultaneously localize targets while doing sensor planning.

We provide a Bayesian formulation that generalizes our previous work [2,6] to include multi-class targets with different mission importance. We propose an objective function that is approximately optimized to yield allocation policies for multiple heterogeneous sensors. The performance of these policies is compared numerically to baseline policies, oracle policies and previously proposed policies. Results indicate adaptive policies can perform significantly better than the baseline policy and come close to the oracle in many cases. Moreover, policies which include mixtures of global uniform sensors and local adaptive sensors (practical and cheap) can perform closely to a single global adaptive sensor (possibly impractical or expensive) in many cases.

The rest of the paper is organized as follows: Section II presents the target and observation models; Section III provides an objective function and approximate optimization methods; Section IV compares performance of policies; and the paper concludes in Section V.

II. Model

A. Signal model

Consider a scene of dimension N, which contain a small number of targets. The i-th location is characterized by its target class $C_i \in \mathcal{C} = \{0, 1, \cdots\}$, where $C_i = 0$ indicates the absence of a target. Define $\{p_c\}_{c\mathcal{C}}$ as the prior distribution of the target class, then the sparse targets assumption can be restated as $p_0 \approx 1$. Correct classification of a location to class c yields a known reward $h(c)$ to the mission (hereafter referred to as the mission importance), and $h(0) = 0$. For example, in a 3-class problem, we might consider $c = 0, 1, 2$ with $h(c) = 0, 1, 10$, respectively, to represent the no-target, low-value target and high value target classes.

Associated with each location i is an amplitude X_i, called the signal and modeled as follows:

$$X_i | C_i = c \sim \left\{ \begin{array}{ll} 0, & c = 0 \\ \text{Normal}(\mu_c, \sigma_0^2), & c \neq 0 \end{array} \right. . \tag{1}$$

In words, conditioned on the event $C_i = c$ for $c > 0$, X_i is zero with probability one if $c = 0$, and is distributed as a Gaussian random variable with known mean μ_c and variance σ_0^2 if $c > 0$. It should be noticed that we assume the variance of non-empty classes are the same σ_0^2, this can be restrictive in some cases. However, we will show that it leads to simplified optimization problem and argue that it can be easily extended.

Given total sensing budget Λ, a sequence of T stages of observations are made with resources $\lambda_i(t)$ that vary with index and stage $t = 0, \cdots, T-1$. The exploration resource is application-dependent and can represent extended effort like observation time, number of samples, or battery life. Given $\lambda_i(t)$, the corresponding observation $y_i(t)$ takes the form:

$$y_i(t) = X_i + \frac{n_i(t)}{\sqrt{\lambda_i(t)}} \tag{2}$$

where $\{n_i(t)\}_{i,t}$ is i.i.d. zero-mean Gaussian noise with variance ν^2. If $\lambda_i(t-1) = 0$, no observation is taken. A key point of this model is that the observation quality increases with sensing effort (i.e., $\lambda_i(t-1)$). The resource allocations are constrained with the following total resource budget:

$$\sum_{t=0}^{T-1} \sum_{i=1}^{N} \lambda_i(t) = \sum_{t=0}^{T-1} \Lambda(t) = \Lambda, \tag{3}$$

where $\Lambda(t)$ are the per-stage budgets. For convenience, we write $\boldsymbol{\lambda}(t) = [\lambda_1(t)\ \ \lambda_2(t) \ldots \lambda_N(t)]^T$, $\boldsymbol{y}(t) = [y_1(t)\ y_2(t) \ldots y_N(t)]^T$ (similarly for other indexed quantities) and define $\boldsymbol{Y}(t) = \{\boldsymbol{y}(1), \boldsymbol{y}(2), \ldots, \boldsymbol{y}(t)\}$. The sequence of effort allocations $\boldsymbol{\lambda} = \{\boldsymbol{\lambda}(t)\}_{t=0}^{T-1}$ is called the allocation policy, where $\boldsymbol{\lambda}(t)$ is a mapping from $\boldsymbol{Y}(t)$ to $[0, \Lambda(t)]^N$.

B. Sensor model

In many cases, the system may have various types of sensors available. These sensors typically vary in sensing agility, range, energy consumption, and cost. We consider three types of sensors: global uniform (GU) sensor, global adaptive (GA) sensor, and local adaptive (LA) sensor. The global adaptive sensor is able to chose any subset of the scene and allocate different amounts of sensing resource to dimensions of the subset. These sensors provide the most agility, but may not exist in reality or are very expensive to build and deploy. The global uniform sensor explores all dimensions, but with equal sensing resource in each dimension, for example, a global camera that can monitor the entire scene. Therefore, it is likely to spend resources in locations without targets and suffer performance degradation compared to the global adaptive sensors. Finally, the local adaptive sensor is limited to exploring only a small number of locations within the scene, albeit with high resolution, e.g. UAVs. These sensors tend to be cheap as compared to global adaptive sensors and may be easily available in practice. Table I and Figure 1 compares these sensor types with regard to their allocations.

TABLE I: Comparison of sensor types and allocations

Sensor Type	Allocation	Adaptivity	Availability
Global Adapt.	$\lambda_i(t) \in [0, \Lambda/T]$	Full	Low
Global Unif.	$\lambda_i(t) = \Lambda/(NT)$	None	High
Local Adapt.	$\lambda_i(t) = k\Lambda/(MT)$, $k = \{0, 1, \ldots, M\}$	Limited	Medium

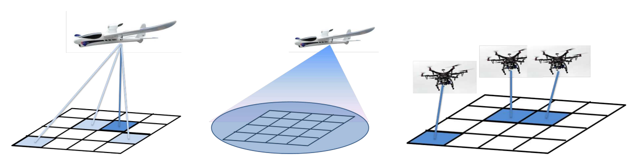

(a) Global adaptive (b) Global uniform (c) Local adaptive

Fig. 1: Comparison of three types of sensors. GA may allocate differently to each location. GU allocates uniformly over the scene. LA allocates discrete units to a subset of the scene.

III. Search Policy

A. Objective function

Define indicator functions $I_i^{(c)}$:

$$I_i^{(c)} = \begin{cases} 1, & C_i = c \\ 0, & C_i \neq c \end{cases}. \tag{4}$$

At stage t, Denote the posterior mean (conditional mean estimator), variance and classification probability as are defined as $\hat{X}_i^{(c)}(t)$, $(\sigma_i^{(c)}(t))^2$ and $p_i^{(c)}(t)$:

$$\hat{X}_i^{(c)}(t) = \mathbb{E}\left[X_i \middle| \boldsymbol{Y}(t), C_i = c\right] \tag{5}$$

$$(\sigma_i^{(c)}(t))^2 = \mathrm{var}\left[X_i \middle| \boldsymbol{Y}(t), C_i = c\right] \tag{6}$$

$$p_i^{(c)}(t) = \Pr(C_i = c | \boldsymbol{Y}(t)) \tag{7}$$

with $\hat{X}_i^{(c)}(0) = \mu_c$, $(\sigma_i^{(c)}(0))^2 = \sigma_c^2$, and $p_i^{(c)}(0) = p_c$.

In this work, we consider a generalization of the cost function in [2] to multiple target classes that accounts for the mission value of each class $h(c)$:

$$J_T(\boldsymbol{\lambda}) = \mathbb{E}\left[\sum_{i=1}^{N} \sum_{c \in \mathcal{C}} h(c) I_i^{(c)} (X_i - \hat{X}_i^{(c)}(T))^2\right], \tag{8}$$

To simplify further, notice that the posterior covariance $(\sigma_i^{(c)}(t))^2$ follows the following update rule:

$$\frac{1}{(\sigma_i^{(c)}(t+1))^2} = \frac{1}{(\sigma_i^{(c)}(t))^2} + \frac{\lambda_i(t)}{\nu^2} = \frac{1}{\sigma_0^2} + \frac{\sum_{t'=1}^{t} \lambda_i(t')}{\nu^2}, \tag{9}$$

which does not depend on class c[1]. Subsequently, we define the class-independent posterior variance as

$$\sigma_i^2(t+1) = \left[\frac{1}{\sigma_i^2(t)} + \frac{\lambda_i(t)}{\nu^2}\right]^{-1} = \left[\frac{1}{\sigma_0^2} + \frac{\sum_{t'=1}^{t} \lambda_i(t')}{\nu^2}\right]^{-1} \tag{10}$$

As in [2], it can be shown that

$$\mathbb{E}_{\boldsymbol{y}(T)}\left[p_i^{(c)}(T) | \boldsymbol{Y}(T-1)\right] = p_i^{(c)}(T-1) \tag{11}$$

Thus, plugging into (8), the cost function can be re-written as

$$\begin{aligned} J_T(\boldsymbol{\lambda}) &= \mathbb{E}\left[\sum_{i=1}^{N} \sum_{c \in \mathcal{C}} h(c) p_i^{(c)}(T) \sigma_i^2(T)\right] \\ &= \nu^2 \mathbb{E}\left[\sum_{i=1}^{N} \sum_{c \in \mathcal{C}} \frac{h(c) p_i^{(c)}(T-1)}{\nu^2/\sigma_i^2(T-1) + \lambda_i(T-1)}\right] \\ &= \nu^2 \mathbb{E}\left[\sum_{i=1}^{N} \frac{z_i(T-1)}{\nu^2/\sigma_i^2(T-1) + \lambda_i(T-1)}\right], \end{aligned} \tag{12}$$

where

$$z_i(T-1) = \sum_{c \in \mathcal{C}} h(c) p_i^{(c)}(T-1) \tag{13}$$

[1] This is a result of assuming that the prior variances for each class are equal, see Sec. II.

B. Global adaptive (GA) policy

As discussed in previous work on similar problems [2], it is possible in principle to use dynamic programming (DP) to optimize $J_T(\boldsymbol{\lambda})$ over $\boldsymbol{\lambda}$ given resource constraints (3). However, for $T > 2$, the exact solution is computationally intractable. Here we present a myopic method which independently defines $\boldsymbol{\lambda}(t)$ for each $t = 0, 1, 2, \ldots, T-1$ as follows:

$$\begin{aligned} \min_{\boldsymbol{\lambda}(t)} \quad & \mathbb{E}\left[\sum_{i=1}^{N} \frac{z_i(t)}{\nu^2/\sigma_i^2(t) + \lambda_i(t)}\right] \\ \text{s.t.} \quad & \sum_{i=1}^{N} \lambda_i(t) = \Lambda(t). \end{aligned} \tag{14}$$

where $\Lambda(t)$ is some fraction of the total budget Λ. Given $\Lambda(t)$, the optimal solution follows [2] and begins by defining π to be an index permutation that sorts $\sqrt{z_i(t)}\sigma_i^2(t)$ in non-increasing order:

$$\sqrt{z_{\pi(1)}(t)}\sigma_{\pi(1)}^2(t) \geq \cdots \geq \sqrt{z_{\pi(N)}(t)}\sigma_{\pi(N)}^2(t). \tag{15}$$

Let $c_i(t) = \nu^2/\sigma_i^2(t)$ and define $g(k)$ to be the monotonically non-decreasing function of $k = 0, \ldots, N$ with $g(0) = 0$,

$$g(k) = \frac{c_{\pi(k+1)}(t)}{\sqrt{z_{\pi(k+1)}(t)}} \sum_{i=1}^{k} \sqrt{z_{\pi(i)}(t)} - \sum_{i=1}^{k} c_{\pi(i)}(t), \tag{16}$$

for $k = 1, \ldots, N-1$, and $g(N) = \infty$. Then the solution to (14) is

$$\lambda_{\pi(i)}^{ga}(t) = \left(\Lambda(t) + \sum_{j=1}^{k^*} c_{\pi(j)}(t)\right) \frac{\sqrt{z_{\pi(i)}(t)}}{\sum_{j=1}^{k^*} \sqrt{z_{\pi(j)}(t)}} - c_{\pi(i)}(t), \tag{17}$$

for $i = 1, \ldots, k^*$ and $\lambda_i^{ga}(t) = 0$ else. The number of nonzero components is determined by the interval $(g(k-1), g(k)]$ to which the budget parameter $\Lambda(t)$ belongs. Since $g(k)$ is monotonic, the mapping from $\Lambda(t)$ to k^* is well-defined. We note that it is possible to optimize over the budget fractions $\Lambda(t)$ as in [2], but for simplicity we consider a fixed budget $\Lambda(t) = \Lambda/T$ for $t = 1, 2, \ldots, T$.

C. Local adaptive (LA) policy

In some cases it may be impossible to deploy a sensor with the agility to assign different resources to every location in the scene at every stage t. Instead we consider the situation where there are M local sensors (e.g. UAVs) that can explore a subset of the locations with a fixed resource amount per sensor. At each stage $t = 0, 1, \ldots, T-1$, define $u_i(t) \in \{0, 1, \ldots, M\}$ as the number of sensors allocated to location i for $i \in \mathcal{X}$, $\sum_{i=1}^{N} u_i(t) = M$, and $\boldsymbol{u}(t) = \{u_i(t)\}_{i=1}^{N}$. Then the LA allocation problem becomes:

$$\begin{aligned} \min_{\boldsymbol{u}(t)} \quad & \mathbb{E}\left[\sum_{i=1}^{N} \frac{z_i(t)}{\nu^2/\sigma_i^2(t) + \lambda_i(t)}\right] \\ \text{s.t.} \quad & \lambda_i^{la}(t) = u_i(t) \frac{\Lambda(t)}{M} \end{aligned} \tag{18}$$

for $t = 0, 1, 2, \ldots, T-1$. The solution is given by a series of M steps, where at each step, we allocate a single sensor to the location that provides the greatest decrease in the cost function. First, set $\hat{u}_i(t) = 0$ for $i = 1, 2, \ldots, N$. Then, define the decrease in cost by allocating a single additional sensor to location i as

$$\Delta_i(t) = \frac{z_i(t)}{\nu^2/\sigma_i^2(t) + \hat{u}_i(t)\Lambda(t)/M} \tag{19}$$

$$- \frac{z_i(t)}{\nu^2/\sigma_i^2(t) + (\hat{u}_i(t)+1)\Lambda(t)/M}. \tag{20}$$

In each of M steps, we set $\hat{u}_{i^*}(t) \leftarrow \hat{u}_{i^*}(t) + 1$, where i^* is the index with the largest $\Delta_i(t)$. Note that multiple sensors are allowed to visit the same location.

The greedy allocation above is optimal for (18) because of the convexity of the myopic cost with respect to each $\lambda_i(t)$. As a consequence, the decreases $\Delta_i(t)$ at a particular location diminish as more sensors are assigned to it, and we can be sure that the assignment in each step yields the largest decrease globally.

D. Global uniform/local adaptive (GU/LA) mixture policy

With no prior information on the location of targets in the scene, the local policy is likely to perform poorly, because it may take a long time to locate targets if only a small number of positions can be queried in each stage. Thus, we consider a third sensing modality where a global sensor is able to gather low resolution information on target location by uniformly observing the scene. Subsequently, local sensors can use this information to measure the likely locations with higher signal quality. In this policy, the optimization is two-fold: (a) the percentage of resources used by the global exploration sensor is optimized; and (b), the locations of the local sensors are optimized. This optimization problem can be written as

$$\min_{T_s, \{\boldsymbol{u}(t)\}_{t=0}^{T-1}} \quad J_T(\boldsymbol{\lambda}^{gula}) \tag{21}$$

$$\text{s.t.} \quad \lambda_i^{gula}(t; \boldsymbol{u}(t), T_s) = \begin{cases} \Lambda(t)/N, & t \leq T_s \\ u_i(t)\Lambda(t)/M, & t > T_s \end{cases}$$

Given T_s, the optimization problem reduces to finding the sensor allocations $\boldsymbol{u}(t)$ for $t = T_s + 1, \ldots, T-1$, which is done myopically using steps in the Local Adaptive policy. Optimization over T_s is done offline (i.e., before any real measurements are taken) by approximating the expectation in $J_T(\boldsymbol{\lambda})$ through Monte Carlo samples. In general, this requires $\mathcal{O}(T)$ Monte Carlo samples to determine the optimal T_s in a T-stage policy.

IV. SIMULATION

This section presents a numerical study for comparing the performance of proposed policies — global adaptive (GA), local adaptive (LA), and global uniform/local adaptive mixture (GU/LA), with two benchmark policies, the global uniform (GU) that allocate resource uniformly everywhere and the oracle policy that knows the target locations as a priori, and a previously proposed policy ARAP [1,2], which is designed to only detect targets and not classify them.

Simulation parameters are given in Table II unless otherwise stated. In this scenario, there are very few targets in the scene (5% on average), which leads to significant performance gaps between the GU policy and adaptive policies. Moreover, the condition that $\mu_3 < \mu_2$ is imposed to account for the fact that high-importance targets are generally harder to detect than low-importance targets. We set $T = 10$ for the GA policy and $T = 30$ for the GU/LA and LA policies since the latter two require additional stages to effectively search the entire scene. The total budget is defined as a function of SNR: $\Lambda = 10^{\mathrm{SNR}/10}N$.

TABLE II: Simulation parameters

Parameter	Value
Number of locations, N	2500
Number of classes, $\lvert\mathcal{C}\rvert$	3
Class probabilities, $\{p_c\}_{c\in\mathcal{C}}$	$\{0.95, 0.049, 0.001\}$
Class importance, $\{h(c)\}_{c\in\mathcal{C}}$	$\{0, 1, 2500\}$
Target prior means, $\{\mu_c\}_{c\in\mathcal{C}}$	$\{0, 3, 1.5\}$
Target prior variances, $\{\sigma_c^2\}_{c\in\mathcal{C}}$	$\{0, 1/16, 1/16\}$
Noise variance, ν^2	1
Number of sensors for LA policy, M	400
Number of sensors for GU/LA policy, M	50

Figs. 2a and 2b explore the sensitivity of the GU/LA policy to the percentage of total resources used by the global sensor, T_s/T. In Fig. 2a, it is seen that there is significant performance degradation when this percentage is close to either extreme. When $T_s = T$ (i.e., the GU policy), the lack of adaptivity leads to inefficient resource allocation. Conversely, when $T_s = 0$ (i.e., the LA policy), the local sensors have difficulty in finding the locations that contain valuable targets. Nevertheless, for any fixed SNR, the gain of the GU/LA policy is relatively flat in a region around the optimal T_s value. Circles indicate the maximum T_s/T within 3 dB of the maximum gain, while diamonds indicate the minimum T_s/T within 3 dB. It is seen that for any SNR, there is a large region where the gain is within 3 dB, which suggests that the GU/LA policy may be robust to small errors in finding the optimal T_s. Fig. 2b shows the optimal T_s values while varying both SNR and p_3.

(a) $J_T(\boldsymbol{\lambda}^u)/J_T(\boldsymbol{\lambda}^{gula})$ (b) Optimal T_s vs. SNR and p_3

Fig. 2: Sensitivity of GU/LA to T_s/T. In (a) For each SNR value, the optimal percentage is given by black squares. Circles and diamonds indicate the maximum/minimum T_s/T, respectively, where the gain is within 3 dB of the maximum. There is significant degradation when $T_s = 0$ (i.e. the LA policy) or $T_s = T$ (i.e., the GU policy). (b) shows the optimal T_s/T for the GU/LA policy while varying both SNR and p_3.

Fig. 3 explores the benefit of including a global sensor by comparing the GU/LA and LA policies. Figs. 3a and 3b show the gains of the LA and GU/LA policies with respect to the GU policy, while varying SNR and the number of local sensors. In this simulation, the optimal T_s for the GU/LA policy is given

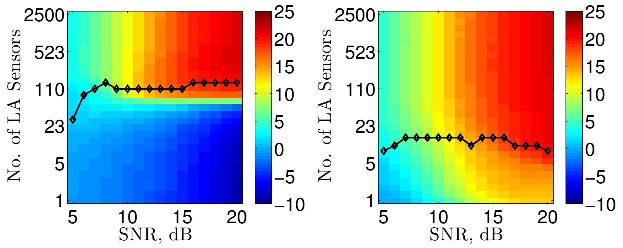

(a) $J_T(\boldsymbol{\lambda}^u)/J_T(\boldsymbol{\lambda}^{la})$ (b) $J_T(\boldsymbol{\lambda}^u)/J_T(\boldsymbol{\lambda}^{gula})$

Fig. 3: Benefit by including a global sensor. (a) and (b) show the gains as a function of SNR and the number of local sensors. For each SNR value, a black diamond indicates the smallest number of sensors needed to achieve gains within 3 dB of the maximum gain. Both policies have very good performance when the number of local sensors is large. However, the LA policy requires at least 100 sensors to acheive within-3dB performance in most cases, and suffers large decreases in performance when this condition is not satisfied. On the other hand, the GU/LA policy requires many fewer sensors (on the order of 10 sensors) to achieve within-3dB performance.

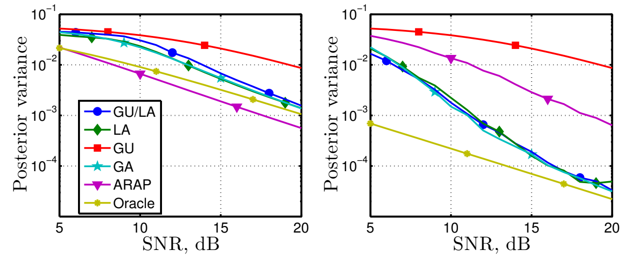

(a) low-importance targets (b) high-importance targets

Fig. 4: Posterior variance of low-importance and high-importance targets. GU performs the worst as it does not adapt resource allocation. The GA, LA, and GU/LA policies perform better than ARAP over high-importance targets, but worse on low-importance targets. Note that the LA policy uses 400 local sensors, while the GU/LA policy only uses 50 sensors in addition to the GU sensor.

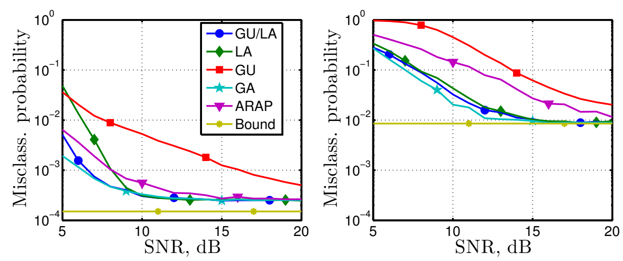

(a) low-importance targets (b) high-importance targets

Fig. 5: Misclassification probability as function of SNR and policy. All adaptive policies perform significantly better than GU and approach the performance of the oracle as SNR gets large. Once again LA, GU/LA, and GA perform better than ARAP in reducing misclassification errors for the high-importance targets.

by the values in Fig. 2a. Black diamonds indicate the minimum number of sensors needed to achieve within 3 dB of the maximum gain at each SNR. Note that the maximum number of sensors is equal to $N = 2500$ for both policies. However, the LA policy requires at least 100 local sensors in almost all cases to attain performance within 3 dB of the maximum, while the GU/LA policy requires an order of magnitude fewer sensors. Furthermore, a phase transition occurs for the LA policy when fewer than 100 sensors are used, wherein the LA policy actually performs worse than the GU policy.

The proposed policies approximately optimize an objective function that combines estimation and classification errors. However, the policies also perform well in reducing each of these errors individually as demonstrated in the next figures. Fig. 4 plots the posterior variance (i.e., expected estimation error) within the low-value and high-value target classes. Fig. 5 shows the misclassification probabilities within each class. Fig. 5 also includes a bound on these probabilities in the asymptotic case where $\Lambda \to \infty$.

It is seen that the GA, GU/LA, and LA policies all perform similarly to the oracle policy in all of these metrics as SNR improves. ARAP, which does not distinguish between high- and low-value targets, performs better for low-value targets and worse for high-value targets. Note that in these plots, the GU/LA policy used only 50 local sensors in comparison to the LA policy which used 400 local sensors. The GA policy has slightly lower misclassification errors than the GU/LA and LA policies for low SNR values (SNR $<$ 10 dB).

In some cases, the ultimate goal of the mission is to use the information obtained by the sensors in order to achieve some objective. For example, this may include dropping payloads for military purposes or for first-aid after natural disasters. Moreover, it may be the case that there are only a limited number of payloads available. While the goal of this paper is not to optimize the expected return on these payloads, the performance of the proposed policies as a function of the number of payloads is easily simulated. Define the expected return at location i as the posterior mean reward at the final stage $z_i(T)$, as defined by (13). Thus payloads are assigned to the p locations with highest expected importance. Define w as a sorting operator such that $z_{w(1)} \geq \cdots \geq z_{w(N)}$. Then the total return with p payloads is

$$P_T(p) = \sum_{i=1}^{p} z_{w(i)}(T), \tag{22}$$

Fig. 6 compares $P_T(p)$ as a function of p and policy with SNR = 8 dB. The GA, GU/LA, and LA policies quickly converge to the oracle policy, followed by ARAP. However the LA policy requires 5 times more local sensors and has larger signal estimation error than the GU/LA policy.

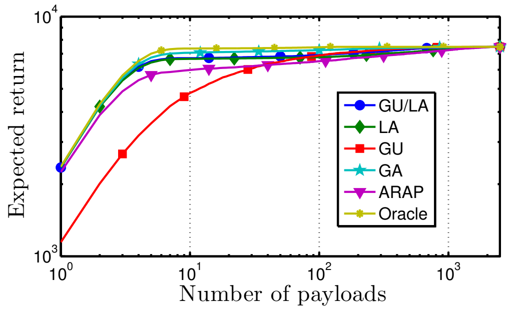

Fig. 6: The expected return given by (22) as function of policy, when there are a limited number of payloads. The GA policy quickly approaches the oracle policy, followed by the GU/LA and LA policies. The ARAP and GU policies have slower convergence to the oracle policy return.

V. Conclusion

This paper generalized previous work on adaptive target search to include value of information where targets have heterogeneous mission importance. New polices are proposed that are able to simultaneously locate, classify and estimate a sparse number of targets embedded in a large space. We also considered heterogeneous sensor models, which include global uniform sensors, global adaptive sensors, and local adaptive sensors. Simulation results show adaptive policies provide significant performance gains over a uniform global allocation policy. Moreover, a policy which combines global and local sensors can yield comparable performance as that with a global adaptive sensor, especially when SNR is low or the number of payloads is small, but the local adaptive and global uniform sensors tend to be cheaper and more practical than a full-scale global adaptive sensor.

Future work will (a) compare and contrast this formulation to one which directly optimizes the expected return rather than our proposed objective function; (b) include time-varying mission-value that will model stale information; and (c) develop performance bounds and theoretical guarantees on performance, such as lower/upper bounds on estimation error.

Acknowledgment

The research in this paper was partially supported by Army Research Office MURI grant number W911NF-11-1-0391.

References

[1] E. Bashan, R. Raich, and A. Hero, "Optimal Two-Stage Search for Sparse Targets Using Convex Criteria," *IEEE Transactions on Signal Processing*, vol. 56, pp. 5389–5402, 2008.

[2] D. Wei and A. Hero, "Multistage adaptive estimation of sparse signals," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 7, no. 5, pp. 783–796, 2013.

[3] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *Information Theory, IEEE Transactions on*, vol. 57, no. 9, pp. 6222–6235, 2011.

[4] A. Averbuch, S. Dekel, and S. Deutsch, "Adaptive compressed image sensing using dictionaries," *SIAM Journal on Imaging Sciences*, vol. 5, no. 1, pp. 57–89, 2012. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/110820579

[5] P. Indyk, E. Price, and D. P. Woodruff, "On the power of adaptivity in sparse recovery," *Foundations of Computer Science, IEEE Annual Symposium on*, vol. 0, pp. 285–294, 2011.

[6] E. Bashan, G. Newstadt, and A. Hero, "Two-stage multiscale search for sparse targets," *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2331–2341, 2011.

[7] A. Tajer, R. Castro, and X. Wang, "Adaptive sensing of congested spectrum bands," *Information Theory, IEEE Transactions on*, vol. 58, no. 9, pp. 6110–6125, 2012.

[8] E. K. Chong, C. M. Kreucher, and A. O. Hero, Iii, "Partially observable markov decision process approximations for adaptive sensing," *Discrete Event Dynamic Systems*, vol. 19, no. 3, pp. 377–422, Sep. 2009. [Online]. Available: http://dx.doi.org/10.1007/s10626-009-0071-x

[9] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli, "Decision-theoretic coordination and control for active multi-camera surveillance in uncertain, partially observable environments," in *Distributed Smart Cameras (ICDSC), 2012 Sixth International Conference on*, 2012, pp. 1–6.

[10] V. Krishnamurthy, "Emission management for low probability intercept sensors in network centric warfare," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 41, no. 1, pp. 133–151, 2005.

[11] J. Capitan, M. Spaan, L. Merino, and A. Ollero, "Decentralized multi-robot cooperation with auctioned pomdps," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, 2012, pp. 3323–3328.

[12] B. Mu, G. Chowdhary, and J. P. How, "Value-of-information aware active task assignment," in *AIAA Guidance Navigation and Control*, Boston, August 2013. [Online]. Available: http://arc.aiaa.org/doi/pdf/10.2514/6.2013-4997

[13] L. Bertuccelli and J. How, "Active exploration in robust unmanned vehicle task assignment," *Journal of Aerospace Computing, Information, and Communication*, vol. 8, pp. 250–268, 2011. [Online]. Available: http://dx.doi.org/10.2514/1.50671