# ADAPTIVE SENSING RESOURCE ALLOCATION OVER MULTIPLE HYPOTHESIS TESTS

*Dennis Wei*

Thomas J. Watson Research Center, IBM Research, Yorktown Heights, NY 10598, USA

**ABSTRACT**

This paper considers multiple binary hypothesis tests with adaptive allocation of sensing resources from a shared budget over a small number of stages. A Bayesian formulation is provided for the multistage allocation problem of minimizing the sum of Bayes risks, which is then recast as a dynamic program. In the single-stage case, the problem is a non-convex optimization, for which an algorithm is presented that ensures a global minimum under a sufficient condition. In the mutistage case, the approximate dynamic programming method of open-loop feedback control is employed. The proposed allocation policies outperform alternative adaptive procedures when the numbers of true null and alternative hypotheses are not too imbalanced. In the case of few alternative hypotheses, the proposed policies are competitive using only a few stages of adaptation. In all cases substantial gains over non-adaptive sensing are observed.

***Index Terms***— Sequential decisions, signal detection, multiple testing, dynamic programming, non-convex optimization.

## 1. INTRODUCTION

This paper is concerned with the problem of multiple binary hypothesis tests under a shared sensing budget. Sensing resources can be allocated adaptively over multiple stages to the hypothesis tests, taking past observations into account. Intuitively, the advantage of adaptive allocation is that resources can be continually shifted from tests where the outcome is more certain to those that are less certain. For example, in wide-area search and surveillance, sensors can be directed to gradually concentrate more time, samples, or energy on regions where target presence is the most uncertain. Other applications include adaptive spectrum sensing [1] and multistage gene association studies [2].

Adaptive and sequential methods for multiple testing have been studied recently in [3–5], particularly in the context of support recovery for sparse signals. These works showed that simple multistage thresholding procedures can asymptotically drive error rates to zero with slower growth in resources compared to non-adaptive procedures; [3] focused on Gaussian observations and false discovery/non-discovery rates, while [4, 5] considered more general likelihoods and the family-wise error rate. The present work differs from [3–5] in three major respects: First, no sparsity assumption is made on the number of alternative (or null) hypotheses that are true. Indeed, significant performance gains are demonstrated even when the hypotheses occur in equal numbers. Second, the number of stages is decoupled from the number of hypothesis tests and is deliberately constrained to be small. It is shown that much of the benefit of adaptation can be realized with only two or three stages. Third, a Bayesian formulation is adopted that allows for composite null and alternative hypotheses given statistical prior knowledge; [3–5] in contrast require a simple null hypothesis but less prior information.

The present paper and [3–5] are related more broadly to the literature on sequential (single) hypothesis testing [6], especially with more than two hypotheses and control over observations [7–11]. However, while it may be possible in principle to apply these methods to multiple tests, performance losses may be expected compared to more specialized methods such as in [3–5] and herein. Moreover, sequential procedures as defined in [7–11] allow an indefinite number of stages at which sensing decisions can be made, in contrast to the approach here where the number of stages is fixed and small and the resource budget is also fixed. In addition, [7–11] consider a finite number of sensing choices of differing quality but equal cost, whereas in this work the sensing control is continuous-valued and quality is a direct function of resource cost.

The statistical model and dynamic programming methods in this paper are similar to those in [12] (except for sparsity). However, the objective of hypothesis testing differs significantly from [12], which focuses on amplitude estimation of sparse signals. This difference has an important consequence for optimization: the Bayes risk adopted here as the performance metric is not a convex function of the resource allocations, unlike the estimation error metrics in [12]. The lack of convexity complicates the resource allocation problem and necessitates an alternative optimization method.

Section 2 presents a Bayesian formulation of multiple binary hypothesis testing with adaptive allocation of sensing resources. Only Gaussian observations are considered in this paper. The multistage allocation problem of minimizing the sum of Bayes risks is then recast as a dynamic program. In Section 3, single-stage and multistage solutions are developed. In the single-stage case, an algorithm is proposed involving parallel single-variable minimizations and an outer search over a Lagrange multiplier. Despite the non-convexity of the Bayes risk objective function, this algorithm can guarantee a global minimum when a sufficient condition is met. In the multistage case, a tractable approximate solution is proposed using open-loop feedback control [13] that improves monotonically as the number of stages increases, similar to [12]. Section 4 presents numerical simulations comparing the proposed allocation policies to [3, 5], demonstrating advantages when the numbers of null and alternative hypotheses are within an order of magnitude of each other. In the highly imbalanced case, the proposed policies remain competitive and achieve most of the gains using two or three stages.

## 2. PROBLEM FORMULATION

We consider $n$ binary hypothesis tests indexed by $i = 1, \ldots, n$. A priori, the $i$th null and alternative hypotheses are true with known probabilities $\mathbb{P}(H_i = 0) = 1 - p_i(0)$ and $\mathbb{P}(H_i = 1) = p_i(0)$, and $H_i$, $H_j$ are statistically independent for $i \neq j$. It is not assumed that $p_i(0) \ll 1$, i.e., the alternative hypothesis is not necessarily rare, unlike in [3–5, 12].

Observations are made in $T$ stages (indexed in parentheses) following a model similar to the one in [12]. The quality of each observation is controlled by the amount of sensing resources allocated to it. Specifically, given resource $u_i(t-1) > 0$, the observation $y_i(t)$

for test $i$ in stage $t$ is conditionally distributed as

$$y_i(t) \mid x_i, u_i(t-1) \sim \mathcal{N}(x_i, \nu^2/u_i(t-1)), \quad t = 1, \ldots, T, \quad (1)$$

so that the precision (inverse variance) increases with $u_i(t-1)$. If $u_i(t-1) = 0$, the observation $y_i(t)$ is not taken. The mean $x_i$ depends on $H_i$ as specified in (3) below. The nominal variance $\nu^2$ is assumed to be known. The observations $y_i(t)$ are independent across tests $i$ and conditionally independent across stages $t$ given $x_i$ and $u_i(t)$, $t = 0, \ldots, T-1$ (but not unconditionally independent).

As an example of the observation model above with $u_i(t-1)$ an integer, (1) results if $u_i(t-1)$ i.i.d. observations, each distributed as $\mathcal{N}(x_i, \nu^2)$, are taken in stage $t$ and $y_i(t)$ is computed as the sample mean. More generally, $u_i(t-1)$ is allowed to take on any non-negative real value to model continuous-valued resources and for mathematical convenience. The resource allocations are constrained by an overall deterministic budget,

$$\sum_{t=0}^{T-1} \sum_{i=1}^{n} u_i(t) = Bn, \quad (2)$$

so that the average budget per test is $B$. This budget constraint couples the hypothesis tests together.

In adaptive sensing, resource allocations can depend causally on all previous observations. Define $\mathbf{y}(t) = (y_1(t), \ldots, y_n(t))$ (similarly for other vectors) and $\mathbf{Y}(t) = \{\mathbf{y}(1), \ldots, \mathbf{y}(t)\}$. Then $u_i(t-1)$ in (1) is in general a function of $\mathbf{Y}(t-1)$. The mappings $\mathbf{Y}(t) \mapsto \mathbf{u}(t)$ are referred to as the resource allocation policy.

The mean parameters $x_i$ in (1) are independent over $i$ and follow Gaussian distributions conditioned on $H_i$,

$$x_i \mid H_i \sim \mathcal{N}\big(\mu_i^{H_i}(0), \sigma_i^{H_i}(0)^2\big), \quad H_i = 0, 1, \quad (3)$$

with known prior parameters $\mu_i^{H_i}(0)$ and $\sigma_i^{H_i}(0)^2$. Hence both null and alternative hypotheses can be composite if $\sigma_i^0(0), \sigma_i^1(0) > 0$, generalizing [3–5, 12]. By interchanging if necessary, it is assumed that $\sigma_i^0(0) \le \sigma_i^1(0)$ without loss of generality.

After all observations have been collected, a decision $\hat{H}_i(T) : \mathbf{Y}(T) \mapsto \{0, 1\}$ is made in each of the hypothesis tests. Performance is measured by the sum of Bayes risks,

$$R = \sum_{i=1}^{n} \mathbb{E}_{H_i, \mathbf{Y}(T)} \Big[ (1 - H_i)\hat{H}_i(T) + cH_i\big(1 - \hat{H}_i(T)\big) \Big], \quad (4)$$

where $\mathbb{E}_{H_i, \mathbf{Y}(T)}$ denotes expectation over $H_i$ and $\mathbf{Y}(T)$, and $c$ is the cost of a Type II error (miss) relative to a Type I error (false alarm). For $c = 1$, (4) is the sum of the probabilities of error in each test, which is a union bound on the family-wise error rate, i.e., the probability of any error. It is also possible to minimize the family-wise error rate directly using an approach similar to the one in this paper, but this is not developed further here.

In summary, the problem is to minimize the Bayes risk sum (4) with respect to the resource allocation policy $\{\mathbf{u}(t)\}$ subject to the total budget constraint (2).

### 2.1. Dynamic programming formulation

Similar to [12], the multistage minimization of the Bayes risk $R$ can be cast as a dynamic program [13], where the state is a belief state summarizing the posterior distributions of $H_i$ and $x_i$ given observations $\mathbf{Y}(t)$. Using [12, Lem. 1] to derive these posterior distributions, it can be shown that the variables $H_i \mid \mathbf{Y}(t)$ remain independent over $i$ with parameters $p_i(t) = \mathbb{P}(H_i = 1 \mid \mathbf{Y}(t))$, and $x_i \mid H_i, \mathbf{Y}(t)$ remain independent Gaussian with means $\mu_i^{H_i}(t) = \mathbb{E}[x_i \mid H_i, \mathbf{Y}(t)]$ and variances $\sigma_i^{H_i}(t)^2 = \operatorname{var}(x_i \mid H_i, \mathbf{Y}(t))$. The posterior parameters evolve according to

$$p_i(t+1) = \frac{p_i(t) f_i^1\big(y_i(t+1); t\big)}{p_i(t) f_i^1\big(y_i(t+1); t\big) + (1 - p_i(t)) f_i^0\big(y_i(t+1); t\big)}, \quad (5a)$$

$$\mu_i^{H_i}(t+1) = \frac{\nu^2 \mu_i^{H_i}(t) + \sigma_i^{H_i}(t)^2 u_i(t) y_i(t+1)}{\nu^2 + \sigma_i^{H_i}(t)^2 u_i(t)}, \quad (5b)$$

$$\sigma_i^{H_i}(t+1)^2 = \frac{\nu^2 \sigma_i^{H_i}(t)^2}{\nu^2 + \sigma_i^{H_i}(t)^2 u_i(t)}, \quad (5c)$$

where in (5a), $f_i^{H_i}(\cdot; t)$ is the probability density function (PDF) of

$$y_i(t+1) \mid H_i, \mathbf{Y}(t) \sim \mathcal{N}\big(\mu_i^{H_i}(t), \sigma_i^{H_i}(t)^2 + \nu^2/u_i(t)\big), \quad (6)$$

and $t = 0$ corresponds to the prior parameters.

Define the belief state as $\boldsymbol{\xi}(t) = (\mathbf{p}(t), \boldsymbol{\mu}(t), \boldsymbol{\sigma}(t)^2, U(t))$, where $\boldsymbol{\mu}(t)$ and $\boldsymbol{\sigma}(t)^2$ include all components indexed by $i$ and $H_i = 0, 1$, and $U(t)$ is the resource budget remaining in stage $t$ with $U(0) = Bn$ initially. This state definition fulfills the requirement for a dynamic program as specified below.

**Proposition 1.** *The Bayes risk sum* (4) *is the expected value of a function only of the state* $\boldsymbol{\xi}(T-1)$ *and control* $\mathbf{u}(T-1)$*,*

$$R = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{Y}(T-1)} \Bigg[ \int_{-\infty}^{\infty} \min\Big\{ \big(1 - p_i(T-1)\big) f_i^0(y; T-1), \; cp_i(T-1) f_i^1(y; T-1) \Big\}\, dy \Bigg], \quad (7)$$

*where the PDFs* $f_i^0(\cdot; T-1)$ *and* $f_i^1(\cdot; T-1)$ *are completely parameterized in* (6) *by* $\boldsymbol{\xi}(T-1)$ *and* $\mathbf{u}(T-1)$*.*

*Proof outline.* The Bayes risk is minimized by the weighted maximum a posteriori (MAP) rule, which can be expressed in terms of $p_i(T)$. Then (7) is obtained by iterating expectations over $y_i(T) \mid \mathbf{Y}(T-1)$ and $\mathbf{Y}(T-1)$, substituting for $p_i(T)$ using (5a), and simplifying. More details are provided in [14]. □

## 3. RESOURCE ALLOCATION POLICIES

This section discusses single-stage and multistage resource allocation policies that minimize the Bayes risk sum (7) under the budget constraint (2). As discussed in Section 3.2, the single-stage policy of Section 3.1 also applies to the last stage of any multistage policy.

### 3.1. Single-stage policy

In the single-stage case $T = 1$, the expectation in (7) is absent and the objective function simplifies. The remaining integral is the Bayes risk of the optimal test between two Gaussian distributions with different means and variances. The Bayes risk can be evaluated by solving a quadratic inequality to determine the decision regions corresponding to the two terms in the minimization in (7), and then computing the Gaussian integrals, i.e., the Type I and Type II error probabilities. These calculations are fairly standard and the details can be found in [14]. Here the integral in (7) is simply denoted as $R_i(u_i; \boldsymbol{\xi}_i)$, where the stage index $T-1$ is suppressed to simplify notation, and $\boldsymbol{\xi}_i$ represents the components of the state with index $i$. The single-stage resource allocation problem is therefore

$$R^*(\boldsymbol{\xi}) = \min_{\mathbf{u}} \sum_{i=1}^{n} R_i(u_i; \boldsymbol{\xi}_i) \quad \text{s.t.} \quad \sum_{i=1}^{n} u_i = U, \;\; u_i \geq 0. \quad (8)$$

Fig. 1(a) shows that the Bayes risk $R_i(u_i; \boldsymbol{\xi}_i)$ is a decreasing but non-convex function of $u_i$ for a particular choice of parameters $\boldsymbol{\xi}_i$. These properties hold in general for other choices of $\boldsymbol{\xi}_i$, implying that (8) is a non-convex optimization problem.

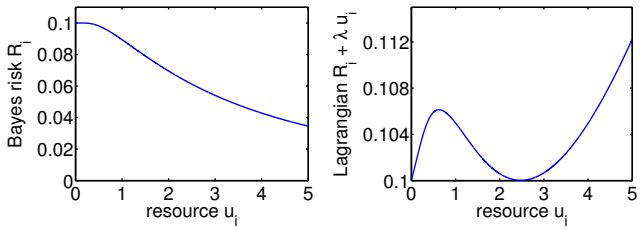


(a) (b)

**Fig. 1**. (a) The Bayes risk $R_i(u_i; \boldsymbol{\xi}_i)$ is a non-convex function of $u_i$. (b) The Lagrangian in (9) can have more than one minimizer.

Despite the absence of convexity, it is still possible in some cases to guarantee a globally optimal solution to (8). We consider minimizing a Lagrangian of (8) in which only the equality constraint is dualized with Lagrange multiplier $\lambda$. The Lagrangian then decouples over $i$. Define the (possibly non-unique) minimizer of each Lagrangian component as

$$u_i(\lambda) \in \arg\min_{u_i \geq 0} \; R_i(u_i; \boldsymbol{\xi}_i) + \lambda u_i. \quad (9)$$

Since $R_i(u_i; \boldsymbol{\xi}_i)$ is bounded from above by $\min\{1 - p_i, cp_i\} = R_i(0; \boldsymbol{\xi}_i)$, a negative value for $\lambda$ in (9) would result in divergence toward infinity. Hence it is sufficient to consider $\lambda \geq 0$. The following result, adapted from [15, Prop. 3.3.4], gives a sufficient condition for $\mathbf{u}(\lambda) = \big(u_1(\lambda), \ldots, u_n(\lambda)\big)$ to be globally optimal for (8).

**Lemma 1.** *If there exists a Lagrange multiplier $\lambda \geq 0$ such that a set of minimizers $\mathbf{u}(\lambda) = \big(u_1(\lambda), \ldots, u_n(\lambda)\big)$ defined by* (9) *is feasible for problem* (8)*, then $\mathbf{u}(\lambda)$ is a global minimum of* (8)*.*

The minimization in (9) also satisfies the monotonicity property below, which confirms the interpretation of $\lambda$ as a penalty parameter. A proof of Lemma 2 is given in [14].

**Lemma 2.** *If $\lambda_1 < \lambda_2$, then $u_i(\lambda_1) \geq u_i(\lambda_2)$ for any minimizers $u_i(\lambda_1)$, $u_i(\lambda_2)$ in* (9)*.*

Based on Lemmas 1 and 2, the following algorithm is proposed to solve (8), consisting of an outer bisection search over $\lambda$ and inner single-variable minimizations (9) to determine $u_i(\lambda)$, $i = 1, \ldots, n$, which can be done in parallel. Lower and upper bounds $\underline{u}_i$ and $\overline{u}_i$ are maintained on each $u_i$, where initially $\underline{u}_i = 0$ and $\overline{u}_i = \infty$. Any algorithm can be used to solve (9) subject to the bounds $\underline{u}_i \leq u_i \leq \overline{u}_i$, for example gradient descent with logarithmically-spaced line search as used to generate the results in Section 4. Let $S(\lambda) = \sum_{i=1}^{n} u_i(\lambda)$. If for a given $\lambda$, the resulting $u_i(\lambda)$ satisfy $S(\lambda) < U$, then $\lambda$ is decreased according to the bisection method, the lower bounds $\underline{u}_i$ are updated to the current solutions $u_i(\lambda)$, exploiting Lemma 2, and (9) is re-solved. Analogous actions are taken if $S(\lambda) > U$. If $S(\lambda) = U$, then by Lemma 1, the algorithm terminates with a globally optimal solution to (8).

For the bisection search over $\lambda$, the initial lower bound is set at 0. The lemma below is used to set the initial upper bound.

**Lemma 3.** *Any minimizer $u_i(\lambda)$ in* (9) *is bounded from above as $u_i(\lambda) < R_i(0; \boldsymbol{\xi}_i)/\lambda = \min\{1 - p_i, cp_i\}/\lambda$.*

*Proof.* Since the Bayes risk $R_i(u_i; \boldsymbol{\xi}_i)$ is positive for finite $u_i$, if $u_i \geq R_i(0; \boldsymbol{\xi}_i)/\lambda$ then $R_i(u_i; \boldsymbol{\xi}_i) + \lambda u_i > R_i(0; \boldsymbol{\xi}_i)$ and $u_i$ cannot be minimal. □

It follows that a sufficient upper bound on $\lambda$ is $\sum_{i=i}^{n} R_i(0; \boldsymbol{\xi}_i)/U$, since any higher value can be seen to result in $S(\lambda) < U$. Lemma 3 is also used to further constrain the inner minimizations over $u_i$ when it gives a tighter upper bound than $\overline{u}_i$.

The above algorithm does not always ensure a global minimum for (8). Specifically, it may not be possible to satisfy the condition in Lemma 1, i.e., there is no $\lambda$ for which $S(\lambda) = U$ to make $\mathbf{u}(\lambda)$ feasible. The problem is illustrated in Fig. 1(b), which shows a value for $\lambda$ such that the Lagrangian in (9) has two separated minimizers. Any change in $\lambda$ would result in either the left or the right minimizer being unique. Hence the function $S(\lambda)$ is discontinuous and the bisection search over $\lambda$ may not converge with $S(\lambda) = U$. For the results in Section 4, cases of non-convergence are addressed simply by rescaling the final solution $\mathbf{u}(\lambda)$ so that it sums to $U$. The loss in optimality appears to be insignificant for large $n$ and can even be bounded analytically, although this is not presented here.

### 3.2. Multistage policies

In a multistage adaptive policy, the last-stage allocation $\mathbf{u}(T-1)$ can depend on all previous observations $\mathbf{Y}(T-1)$. In other words, $\mathbf{u}(T-1)$ is determined after conditioning on $\mathbf{Y}(T-1)$, which again removes the expectation from (7). Therefore the last-stage allocation problem reduces to the single-stage case (8).

For a two-stage policy, it remains to determine the first-stage allocation $\mathbf{u}(0)$. This is done recursively by solving

$$\min_{\mathbf{u}(0)} \; \mathbb{E}_{\mathbf{y}(1)} \left[ R^*(\boldsymbol{\xi}(1)) \mid \boldsymbol{\xi}(0), \mathbf{u}(0) \right] \quad \text{s.t.} \quad \sum_{i=1}^{n} u_i(0) \leq U(0), \quad (10)$$

where $R^*(\boldsymbol{\xi}(1))$ is defined by (8) as the optimal cost of the second stage, and the distribution of $\mathbf{y}(1)$ is parameterized by $\boldsymbol{\xi}(0)$ and $\mathbf{u}(0)$ (see (6)). In the case of priors that are homogeneous over $i$, i.e., $p_i(0)$, $\mu_i^{H_i}(0)$, $\sigma_i^{H_i}(0)^2$ do not depend on $i$ (but can depend on $H_i$), then the first-stage allocation is also homogeneous by symmetry, $u_i(0) = u(0)$, and (10) becomes a scalar minimization with respect to $u(0) \in [0, U(0)/n]$. This minimization is performed offline using Monte Carlo samples of $\mathbf{y}(1)$ to approximate the expectation in (10) and the algorithm in Section 3.1 to approximate $R^*(\boldsymbol{\xi}(1))$.

For an inhomogeneous prior or more than two stages, an open-loop feedback control (OLFC) policy [13] is employed, similar to [12]. Although the derivation of the policy is somewhat more involved than in [12] and is given in [14], the end result is analogous. Specifically, in stage $t$ and conditioned on available observations $\mathbf{Y}(t)$ through the state $\boldsymbol{\xi}(t)$, the single-stage allocation problem (8) with $\boldsymbol{\xi} = \boldsymbol{\xi}(t)$, $U = U(t)$ is solved, resulting in $\mathbf{u}^*$. Then $\mathbf{u}^*$ is scaled by $\beta(t) \in [0, 1]$ to yield $\mathbf{u}(t) = \beta(t)\mathbf{u}^*$, thus conserving some of the resource budget for future stages. The multipliers $\beta(t)$ are optimized offline using Monte Carlo simulation as described in [12]. As shown in [12, Prop. 2], this optimization of $\beta(t)$ ensures that the resulting policies improve monotonically with the number of stages $T$.

## 4. NUMERICAL RESULTS

The multistage resource allocation policies described in Section 3 are numerically compared to the distilled sensing (DS) [3] and sequential thresholding (ST) [5] procedures, as well as to a single-stage non-adaptive baseline policy (NA). For the results presented below, the number of hypothesis tests $n$ is $10^4$ and a homogeneous prior is used: $p_i(0) = p(0)$, $\mu_i^0(0) = 0$, $\mu_i^1(0) = 1$, $\sigma_i^0(0)^2 = 0$, and $\sigma_i^1(0)^2 = 1/16$ for all $i$. Observations are simulated according to (1) and (3). The observation noise parameter $\nu^2$ is normalized to 1 and the average budget per test $B$ is varied. Since $\nu^2$ and $u_i(t)$ always appear in the same ratio as in (1), an equivalent alternative would be to fix $B$ and vary $\nu^2$ instead. The performance metric is (4) with $c = 1$, i.e., it is the expected number of errors of either type.

The number of stages in the proposed OLFC policies is limited between 2 and 4. In all cases, the first-stage allocation $\mathbf{u}(0)$ is uniform because of the homogeneous prior. For $T = 2$, Fig. 2(a) shows the first-stage budget fraction $u(0)$ that results from the offline optimization (10) for different values of $p(0)$ and $B$. For $p(0) = 0.5$, $u(0)$ is relatively constant as a function of $B$, while for small $p(0)$, $u(0)$ increases up until the high-budget regime $B > 10^2$. Performance is not too sensitive to the exact value of $u(0)$ since the objective function in (10) tends to be relatively flat away from the extremes $u(0) = 0$ and $u(0) = 1$. Fig. 2(b) plots $u(0)$ for $T = 3$.

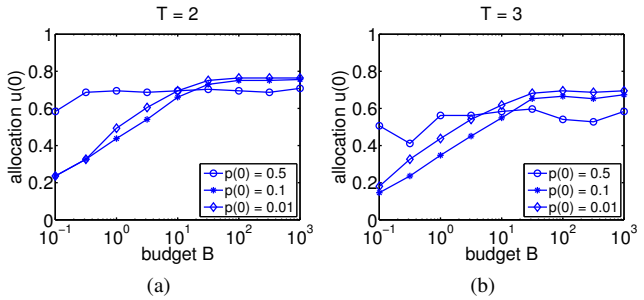


**Fig. 2**. First-stage allocation $u_i(0) = u(0)$ in the proposed 2-stage (a) and 3-stage (b) policies as a function of the mean proportion $p(0)$ of alternative hypotheses and the resource budget per test $B$.

For DS and ST, while [3, 5] prescribe values for $T$ as functions of $n$, in these experiments all $T \in \{2, \ldots, 12\}$ are tested and results for the best $T$ are shown. A similar optimization is performed over the parameter $\rho \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$ in [5]. The budget allocations over stages follow [3, eq. (4),(5)] and [5, eq. (14)] respectively, except in the last stage of ST where the remaining budget is used up entirely. Two versions of DS and ST are implemented: the versions originally proposed in [3, 5] that use only the last stage of observations to make decisions, and Bayesian versions (DSB, STB), not proposed in [3, 5], in which the allocations $\mathbf{u}(t)$ are specified by [3,5] but inference is done through the posterior update equations (5), thus incorporating all stages of observations. As seen below, the Bayesian versions perform considerably better.

The performance of the policies is compared in Fig. 3. For equiprobable hypotheses, $p(0) = 0.5$, the proposed 2-stage policy achieves significant reductions in error (up to a factor of 5) relative to the baseline NA policy, while the 3-stage OLFC policy yields further improvement. Since DS(B) and ST(B) are not designed for this non-sparse scenario, they perform less well, in some cases worse than NA. For $p(0) = 0.1$, the 3-stage OLFC policy essentially dominates the other policies, and at moderate to large resource levels in Fig. 3(d), it is joined by the 2-stage OLFC policy. For $p(0) = 0.01$ and low resources in Fig. 3(e), DSB and STB have slightly lower error rates than the 4-stage OLFC policy, while for higher resources in Fig. 3(f), the opposite is true. Moreover, the 2-stage OLFC policy attains most of the gains of these best-performing policies that use more stages. In particular, the optimized DSB and STB policies shown in Fig. 3 use at least 8 and 6 stages respectively for $B \leq 1$.

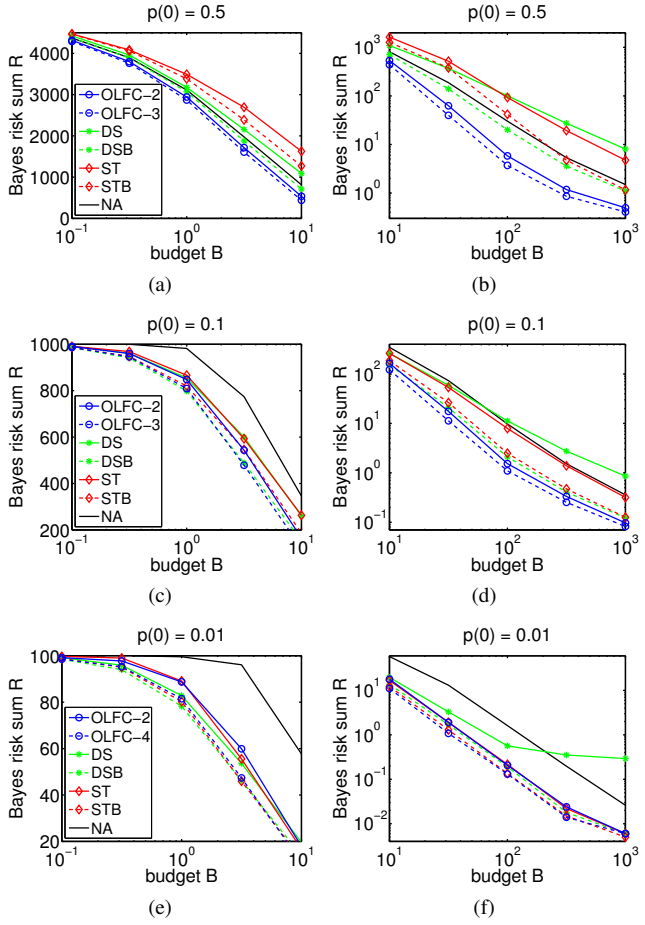


**Fig. 3**. Expected number of errors ((4) with $c = 1$) resulting from the proposed open-loop feedback control policies with $T$ stages (OLFC-$T$), original and Bayesian versions of distilled sensing (DS, DSB) and sequential thresholding (ST, STB), and a non-adaptive baseline (NA). The legends in (a), (c), (e) also apply to (b), (d), (f) respectively. For $p(0) = 0.5, 0.1$ in (a)–(d), OLFC-2 and/or OLFC-3 outperform the alternative methods across budget levels. For $p(0) = 0.01$ in (e)(f), OLFC-4 is competitive with DSB and STB, while OLFC-2 achieves most of the gains using only 2 stages.

## 5. CONCLUSION

This paper has explored the benefits of adaptive sensing for multiple binary hypothesis testing, notably in the regimes of balanced null and alternative hypotheses and few allocation stages. Future work includes generalizations to non-Gaussian observations, refinements of both the single-stage optimization and multistage dynamic programming procedures, and theoretical analysis that aims especially to understand the gains in the non-sparse setting and at moderate, non-asymptotic resource levels.

## 6. REFERENCES


[1] A. Tajer, R. M. Castro, and X. Wang, "Adaptive sensing of congested spectrum bands," *IEEE Trans. Inf. Theory*, vol. 58, no. 9, pp. 6110–6125, Sept. 2012.

[2] S. Zehetmayer, P. Bauer, and M. Posch, "Optimized multi-stage designs controlling the false discovery or the family-wise error rate," *Statist. Med.*, vol. 27, pp. 4145–4160, 2008.

[3] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *IEEE Trans. Inf. Theory*, vol. 57, pp. 6222–6235, Sept. 2011.

[4] M. Malloy and R. Nowak, "On the limits of sequential testing in high dimensions," in *Conf. Rec. Asilomar Conf. Signals Syst. Comput.*, Nov. 2011, pp. 1245–1249.

[5] M. L. Malloy and R. D. Nowak, "Sequential testing for sparse recovery," arXiv:1212.1801v1, Dec. 2012.

[6] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Statist.*, vol. 19, no. 3, pp. 326–339, 1948.

[7] H. Chernoff, "Sequential design of experiments," *Ann. Math. Statist.*, vol. 30, pp. 755–770, 1959.

[8] S. A. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments: Part I—Theory," Tech. Rep. 55, Dept. of Statistics, Stanford Univ., 1960.

[9] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Ann. Statist.*, vol. 41, no. 6, pp. 2703–2738, 2013.

[10] M. Naghshvar and T. Javidi, "Sequentiality and adaptivity gains in active hypothesis testing," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 768–782, Oct. 2013.

[11] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.

[12] D. Wei and A. O. Hero, "Multistage adaptive estimation of sparse signals," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 783–796, Oct. 2013.

[13] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1, Athena Scientific, Nashua, NH, 3rd edition, 2005.

[14] D. Wei, "Adaptive sensing resource allocation over multiple hypothesis tests," arXiv:1410.0950, Oct. 2014.

[15] D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont, MA, 1999.