

# Beyond Binomial and Negative Binomial: Adaptation in Bernoulli Parameter Estimation

Safa C. Medin, *Student Member, IEEE*, John Murray-Bruce, *Member, IEEE*, David Castañón, *Member, IEEE*, and Vivek K Goyal, *Fellow, IEEE*

**Abstract**—Estimating the parameter of a Bernoulli process arises in many applications, including photon-efficient active imaging where each illumination period is regarded as a single Bernoulli trial. Motivated by acquisition efficiency when multiple Bernoulli processes (e.g., multiple pixels) are of interest, we formulate the allocation of trials under a constraint on the mean as an optimal resource allocation problem. An oracle-aided trial allocation demonstrates that there can be a significant advantage from varying the allocation for different processes and inspires the introduction of a simple trial allocation gain quantity. Motivated by achieving this gain without an oracle, we present a trellis-based framework for representing and optimizing stopping rules. Considering the convenient case of Beta priors, three implementable stopping rules with similar performances are explored, and the simplest of these is shown to asymptotically achieve the oracle-aided trial allocation. These approaches are further extended to estimating functions of a Bernoulli parameter. In simulations inspired by realistic active imaging scenarios, we demonstrate significant mean-squared error improvements up to 4.36 dB for the estimation of  $p$  and up to 1.86 dB for the estimation of  $\log p$ .

**Index Terms**—adaptive sensing, Bernoulli processes, beta distribution, coding gain, computational imaging, conjugate prior, dynamic programming, greedy algorithm, lidar intensity, low-light imaging, photon counting, total-variation regularization

## I. INTRODUCTION

Estimating the parameter of a Bernoulli process is a fundamental problem in statistics and signal processing. From the binary-valued outcomes of independent and identically distributed (i.i.d.) trials (generically failure (0) or success (1)), we wish to estimate the probability  $p$  of success. Among myriad applications, our primary interest is raster-scanned active imaging in which a scene patch is periodically illuminated with a pulse, and each illumination period (Bernoulli trial) either has a photon-detection event (success) or not (failure) [1]. The probability  $p$  of a photon-detection event has a monotonic relationship with the reflectivity of the scene patch, and a monotonic function of an estimate of  $p$  becomes the corresponding image pixel value. For efficiency in acquisition time or illumination energy, we are motivated to form the image from a small number of illumination pulses, under conditions where  $p$  is small.<sup>1</sup> Other types of raster-scanned imaging

This material is based upon work supported in part by the US National Science Foundation under Grant No. 1422034 and Grant No. 1815896, and by the DARPA REVEAL program under Contract No. HR0011-16-C-0030.

The authors are with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA.

<sup>1</sup>For applications using time-correlated single photon counting driven by a detector with dead time, such as a single-photon avalanche diode (SPAD), it is recommended to keep  $p$  below 0.05 to avoid time skew and missed detections [2].

can be modeled similarly assuming that the dwell time is an integer multiple of some base time interval, during which the observations are binary.

Conventional methods are not adaptive. With a fixed number of trials  $n$ , the number of successes  $K$  is a binomial random variable, and the maximum likelihood (ML) estimate of  $p$  is  $K/n$ . Though less common in active imaging, a well-known alternative in the statistics literature is to fix a number of successes  $\ell$ . Repeating trials until success  $\ell$  occurs results in a random number of trials  $M$  that is a negative binomial random variable,<sup>2</sup> and the ML estimate of  $p$  is  $\ell/M$ . While there may seem to be nothing to design here, a constraint on the mean number of trials opens up possibilities for data acquisition that results in neither binomial nor negative binomial distributions. The mean may be over a multiplicity of (non-random) Bernoulli process parameters to estimate (such as in active imaging with one parameter per pixel) or over a prior for a single Bernoulli parameter. The two cases are formally linked through the relative frequency interpretation of probability, with the empirical distribution of the multiplicity of deterministic parameters in the former case playing the role of the prior distribution in the latter case [3]. For multiple deterministic parameters, we have a resource allocation problem reminiscent of bit allocation in transform coding [4], [5]. As we will demonstrate, in an oracle-aided setting, trials can be allocated to maximize a *trial allocation gain* that is analogous to the coding gain of transform coding. For a single random parameter, a simple and implementable approach – not requiring an oracle – asymptotically achieves the optimal trial allocation gain and may perform better than the oracle-aided method for moderate numbers of trials.

The focus of this paper is on allocating trials in the estimation of a single random parameter through the design of a stopping rule. A stopping rule may – implicitly and stochastically – allocate trials differently for different values of  $p$ , even though  $p$  is not known *a priori*. We show that any optimal stopping rule can be described by a trellis rather than a more complicated graph, and greedy construction of the trellis is very nearly optimal. For a rectangular array of Bernoulli processes representing a scene in an imaging problem, applying a good stopping rule allocates more trials to the pixels where they provide the most benefit. The final image formation may include a method such as total variation (TV) regularization for exploiting spatial correlations among

<sup>2</sup>Note that the negative binomial distribution is defined inconsistently in the literature, with sometimes the number of failures being fixed rather than the number of successes (reversing the roles of  $p$  and  $1 - p$ ).

neighbors. Regularized image formation makes it more difficult to optimize the acquisition, but it does not invalidate the advantage from adaptive acquisition. In simulations with parameters realistic for active optical imaging, we demonstrate a reduction in mean-squared error (MSE) by a factor of up to 2.73 (4.36 dB) in comparison to the same regularized reconstruction approach applied without adaptation in numbers of trials. Such gains vary based on image content, and gains without regularization are predictable from the trial allocation gain formulation.

### A. Related Work

1) *Statistics Literature:* In statistics, forming a parameter estimate from a number of i.i.d. observations that is dependent on the observations themselves is called *sequential estimation* [6]. Early interest in sequential estimation of a Bernoulli process parameter was inspired by the high relative error of deterministically stopping after  $n$  trials when  $p$  is small. Specifically, the standard error of the ML estimate is  $\sqrt{p(1-p)/n}$ , which for small  $p$  is unfavorable compared to anything proportional to  $p$ . This shortcoming manifests, for example, in requiring large  $n$  to distinguish between two small possible values for  $p$ .

Haldane [7] observed that if one stops after  $\ell$  successes, the (random) number of trials  $M$  is informative about  $p$ . Specifically,  $(\ell - 1)/(M - 1)$  is an unbiased estimate of  $p$  (provided  $\ell \geq 2$ ), and its standard error is proportional to  $p$  (provided  $\ell \geq 3$ ). (The ML estimate  $\ell/M$  is not unbiased, though  $M/\ell$  is an unbiased estimate of  $1/p$ .) Tweedie [8] suggested to call this inverse binomial sampling, but the resulting random variable is now commonly known as negative binomial or Pascal distributed. More recent works have focused on non-MSE performance metrics [9], [10], estimation of functions of  $p$  [11], estimation from imperfect observations [12], and composite hypothesis testing [13].

2) *Photon-Efficient Imaging and Variable Dwell Time:* First-photon imaging [14] introduced sequential estimation to active imaging. This method uses the number of illumination pulses until the first photon is detected to reveal information about reflectivity, setting  $\ell = 1$  in the concept of Haldane [7] and thus using geometric sampling as a special case of negative binomial sampling. A censoring method is used to approximately separate signal and background detections, and spatial correlations are used to regularize the estimation of the full scene reflectivity image, resulting in good performance from only 1 detected photon per pixel, even when half of the detected photons are attributable to uninformative ambient light. Subsequent work with binomial sampling (and otherwise identical experimental conditions) resulted in similar performance [1], and greatly increasing robustness to ambient light is largely attributable to improving the censoring step [15]. These works leave questions on the importance of negative binomial sampling to first-photon imaging unanswered; comparing first-photon imaging to photon-efficient methods with deterministic dwell time [1], [15]–[25] was an initial inspiration for the present work.

While recent works have exploited the first-photon idea in imaging techniques such as ghost imaging [26], [27] and

x-ray tomography [28], previous uses of variable dwell time are not closely connected to sequential estimation or the result of optimized resource allocation. For example, in lidar, varying dwell time to maintain approximately constant signal strength despite varying effective reflectivity (including greater radial fall-off for more distant scene patches) dates back to at least the 1970s [29]. He *et al.* [30] closely follow the technique of [1], including its background censoring, and vary the dwell time to keep the number of photon detections after censoring (i.e., photon detections attributed to signal rather than background) at each pixel approximately constant. In scanning electron microscopy, Dahmen *et al.* [31] increase dwell time where a measure of image detail is large. To the best of our knowledge, no previous paper has formally optimized dwell time under a Bernoulli process measurement model.

### B. Main Contributions and Preview of Results

1) *Framework:* This work discusses a novel framework for depicting and understanding stopping rules for sequential estimation of Bernoulli parameters under number of trials constraints (Section III). In this framework, first presented in [32], each Bernoulli trial corresponds to a transition in a trellis in which each node is identified by the number of trials and number of successes; it is easily shown that distinct paths to reach a given node need not be distinguished. A stopping rule is the assignment of probabilities of stopping to each node in the trellis (see Figs. 4–6). By construction, a stopping rule defined in this way is implementable because it does not depend on knowledge of  $p$  or non-causally on the Bernoulli process. This framework applies equally well under any prior for  $p$ .

2) *Stopping Rule Design:* Simple stopping rules lead to binomial (Fig. 5(a)) and negative binomial (Fig. 5(b)) sampling. Specializing to the Beta family of priors, which is both convenient and conventional because it is the conjugate prior for the relevant observation distributions, methods to optimize the stopping rule are presented in decreasing order of computational complexity: dynamic programming (Section IV-A), offline greedy design (Section IV-B), and online threshold-based termination (Section IV-C) first introduced in [32]. Empirically, all three methods, including the online method requiring no storage of a precomputed stopping rule, provide very similar performance. Thus, the easily implementable online method provides very nearly optimal performance.

3) *Analysis in Oracle-Aided Setting:* This paper introduces the concept of oracle-aided trial allocation whereby processes with different parameters are allocated different fractions of an overall trial budget (Section II-A). This yields a readily-computed *trial allocation gain* that can be arbitrarily large, though it is generally modest (Section II-B). Furthermore, we show that under any Beta prior the threshold-based stopping asymptotically allocates trials identically to the oracle-based optimal (Section IV-E).

4) *Evaluation:* In simulations inspired by realistic active imaging scenarios, an MSE improvement factor of up to 4.36 dB is demonstrated where spatial correlations are

exploited through total variation (TV) regularization (Section VI-A). Without TV regularization, achieved gains are close to the values predicted by the theoretical trial allocation gain. For example, the theoretical trial allocation gain is 2.29 dB, and a gain of 2.27 dB is realized for the Shepp-Logan phantom.

5) *Estimating functions of  $p$* : Inspired by applications where estimating functions of  $p$  is of interest [11], online threshold-based termination is also extended to estimating  $\log p$  from Bernoulli observations (Section V). Experimental results without TV regularization demonstrate improvements of up to 1.86 dB using the threshold-based stopping rule versus the conventional binomial sampling (Section VI-B).

## II. TRIAL ALLOCATION ACROSS MULTIPLE PROCESSES

Consider the estimation of the parameters  $\{p_i\}_{i=1}^r$  of a finite number  $r$  of Bernoulli processes with binomial sampling of each process. When  $m_i$  trials of process  $i$  are observed, the MSE of the ML estimate of  $p_i$  is  $p_i(1-p_i)/m_i$ . Suppose that we are interested in making the average of the MSEs,

$$\frac{1}{r} \sum_{i=1}^r \frac{p_i(1-p_i)}{m_i},$$

small under a constraint on the average of the numbers of trials  $(1/r) \sum_{i=1}^r m_i \leq \eta$ .

Since  $p_i(1-p_i)$  varies over  $[0, 1/4]$  for  $p_i \in [0, 1]$ , there can be an advantage to varying the  $m_i$  values. However, that allocation of trials depends on parameters that are to be estimated. In this section, we suspend the need for implementability and instead study the optimal trial allocation as if the parameters were known. This provides a benchmark for the implementable methods developed in the remainder of the paper, with  $\{p_i\}_{i=1}^r$  playing the role of a discrete prior on  $p$ . We also consider  $r \rightarrow \infty$  to reach a distributional limit.

### A. Oracle-Aided Optimal Allocation

In optimizations such as

$$\min_{m_i, i=1,2,\dots,r} \sum_{i=1}^r \frac{p_i(1-p_i)}{m_i} \quad \text{s.t.} \quad \sum_{i=1}^r m_i \leq r\eta, \quad (1)$$

ignoring that each  $m_i$  should be a positive integer, each MSE vs. number of trials trade-off should be at the same slope, else it would be advantageous to shift trial resources to the process for which the benefit (MSE reduction) per trial is largest. This is formalized using the method of Lagrange multipliers. The resulting optimal allocation is

$$m_i^* = r\eta \frac{\sqrt{p_i(1-p_i)}}{\sum_{j=1}^r \sqrt{p_j(1-p_j)}}, \quad i = 1, 2, \dots, r. \quad (2)$$

Since each process has a fixed number of trials  $m_i^*$ , independent of the experimental outcome of each trial, we call using these numbers of trials *oracle-aided binomial sampling*.

*Example 1 (Oracle-aided allocations):*

- (a) Let  $p_1 = \varepsilon$  and  $p_2 = 1/2$ . Then the fractional oracle-aided allocations are

$$\frac{m_1^*}{2\eta} = \frac{\sqrt{\varepsilon(1-\varepsilon)}}{\sqrt{\varepsilon(1-\varepsilon)} + 1/2}, \quad \frac{m_2^*}{2\eta} = \frac{1/2}{\sqrt{\varepsilon(1-\varepsilon)} + 1/2}.$$




Fig. 1. Oracle-aided optimal allocations in Examples 1(a) and 1(c).

These are plotted as functions of  $\varepsilon$  in Fig. 1(a).

- (b) Let  $p_1 = p_2 = \dots = p_{r-1} = \varepsilon$  and  $p_r = 1/2$ . Then the fractional oracle-aided allocations are

$$\frac{m_i^*}{r\eta} = \frac{\sqrt{\varepsilon(1-\varepsilon)}}{(r-1)\sqrt{\varepsilon(1-\varepsilon)} + 1/2}, \quad i=1, 2, \dots, r-1,$$

$$\frac{m_r^*}{r\eta} = \frac{1/2}{(r-1)\sqrt{\varepsilon(1-\varepsilon)} + 1/2}.$$

- (c) Let  $p_i = (2i-1)/(2r)$ ,  $i = 1, 2, \dots, r$ . The fractional oracle-aided allocations,  $m_i^*/(r\eta)$  are plotted for  $r = 20$  in Fig. 1(b).

### B. Trial Allocation Gain

Using the oracle-aided allocations (2) reduces the average MSE relative to a constant allocation  $m_1 = m_2 = \dots = m_r = \eta$ . The constant allocation results in the average MSE

$$\frac{1}{r} \sum_{i=1}^r \frac{p_i(1-p_i)}{\eta}, \quad (3)$$

whereas using (2) yields

$$\frac{1}{r} \sum_{i=1}^r \frac{p_i(1-p_i)}{m_i^*} = \frac{1}{r} \sum_{i=1}^r \frac{p_i(1-p_i)}{\sqrt{p_i(1-p_i)} r\eta} \frac{1}{\sum_{j=1}^r \sqrt{p_j(1-p_j)}} = \frac{1}{r^2\eta} \left( \sum_{j=1}^r \sqrt{p_j(1-p_j)} \right)^2. \quad (4)$$

We define the ratio of (3) and (4) as the *trial allocation gain*:

$$\gamma_{\text{alloc}} = \frac{r \sum_{i=1}^r p_i(1-p_i)}{\left( \sum_{j=1}^r \sqrt{p_j(1-p_j)} \right)^2}. \quad (5)$$

Trial allocation gain is reminiscent of the coding gain in transform coding [4], [5].

*Example 2 (Trial allocation gains):*

- (a) For the parameters in Example 1(a), the trial allocation gain is plotted as a function of  $\varepsilon$  in Fig. 2. Notice that in the limit of  $\varepsilon \rightarrow 0$ , all the trials are allocated to the nontrivial Bernoulli process, doubling its number of trials, which halves the average MSE. Thus  $\gamma_{\text{alloc}} \rightarrow 2$ .
- (b) For the parameters in Example 1(b),  $\lim_{\varepsilon \rightarrow 0} \gamma_{\text{alloc}} = r$ .
- (c) For the parameters in Example 1(c), the trial allocation gain is plotted as a function of  $r$  in Fig. 2(b).
- (d) Fig. 3(a) shows the “Modified Shepp-Logan phantom” provided by the Matlab `phantom` command, at size  $100 \times 100$  and scaled to  $[0.001, 0.101]$ . Fig. 3(b) shows a histogram of the  $10^4$  intensity values of the phantom. Evaluating (5) gives 1.6944, or 2.29 dB.




Fig. 2. Trial allocation gains in Examples 2(a) and 2(c).




Fig. 3. Phantom of size  $100 \times 100$  with trial allocation gain of 1.686 (Example 2(d)).

One can show that  $\gamma_{\text{alloc}} \in [1, r]$ . The upper bound is illustrated in part (b) of the example.

### C. Distributional Limit

Suppose now that the Bernoulli process parameter is modeled with random variable  $P$  and the number of trials  $M$  is to be assigned by an oracle (i.e., it is allowed to depend on the realization of  $P$ ) to minimize the MSE of the ML estimate of  $P$  under the constraint  $\mathbb{E}[M] \leq \eta$ . By analogy to the computations giving (2) – or formally taking a limit of  $r \rightarrow \infty$  with the empirical distribution of  $\{p_i\}_{i=1}^r$  converging to the distribution of  $P$  – the number of trials should be assigned based on how large  $\sqrt{P(1-P)}$  is relative to  $\mathbb{E}[\sqrt{P(1-P)}]$ :

$$M = \eta \frac{\sqrt{P(1-P)}}{\mathbb{E}[\sqrt{P(1-P)}]}. \quad (6)$$

The resulting trial allocation gain is

$$\gamma_{\text{alloc}} = \frac{\mathbb{E}[P(1-P)]}{\left(\mathbb{E}[\sqrt{P(1-P)}]\right)^2}. \quad (7)$$

This can also be written as

$$\gamma_{\text{alloc}} = 1 + \frac{\text{var}(V)}{(\mathbb{E}[V])^2}, \quad (8)$$

where  $V = \sqrt{P(1-P)}$ . It follows that  $\gamma_{\text{alloc}} \geq 1$ , with equality if and only if the random variable  $\sqrt{P(1-P)}$  has zero variance.

*Example 3 (Trial allocation gains – random parameter):*

- (a) Let  $P$  have the continuous uniform distribution on  $[0, 1]$ . Then evaluating (7) gives  $\gamma_{\text{alloc}} = 32/(3\pi^2) \approx 1.0808$ . This value is the asymptote in Fig. 2(b).
- (b) Let  $P$  take two values:  $\frac{1}{2}$  with probability  $\delta$  and 0 with probability  $1 - \delta$ . Then  $\mathbb{E}[P(1-P)] = \delta/4$  and  $\mathbb{E}[\sqrt{P(1-P)}] = \delta/2$ . Substituting in (7) gives  $\gamma_{\text{alloc}} = 1/\delta$ . We can interpret this with relative frequencies:

Since  $p = 0$  requires no trials, fraction  $\delta$  of the time,  $p = 1/2$  will occur and should be allocated  $1/\delta$  times the mean number of trials.

- (c) When  $p \ll 1$  holds,  $p(1-p) \approx p$ . Therefore, (8) becomes approximately invariant to rescaling. For example, rescaling the phantom in Example 2(d) by a factor of 2 to  $[0.002, 0.202]$  gives  $\gamma_{\text{alloc}} \approx 1.6633$ , and by a factor of  $\frac{1}{2}$  to  $[0.0005, 0.0505]$  gives  $\gamma_{\text{alloc}} \approx 1.7096$ ; these are small changes from the value in Example 2(d).

The first two parts of the example show that though an allocation gain may typically be modest, it may also be arbitrarily large. The third part shows that allocation gain is approximately dependent on the coefficient of variation of the Bernoulli parameter, provided that the parameter is known to be small.

Having established that varying the numbers of trials can be beneficial, we now turn our attention to methods that do not depend on an oracle. We will compare to the oracle-aided allocations in certain asymptotic settings.

### III. OBSERVATION OF A SINGLE BERNOULLI PROCESS

Let  $\{X_n : n = 1, 2, \dots\}$  be a Bernoulli process with an unknown random parameter  $p$ , and let  $\eta \in \mathbb{R}^+$  be a *trial budget*. A *stopping rule* consists of a sequence of *continuation probability functions*

$$\pi_n : \{0, 1\}^n \rightarrow [0, 1], \quad n = 0, 1, \dots, \quad (9)$$

that give the probability of continuing observations after trial  $n$  – based on a biased coin flip independent of the Bernoulli process – as a function of  $(X_1, X_2, \dots, X_n)$ . The result is a random number of observed trials  $N$ .<sup>3</sup> The stopping rule is said to satisfy the trial budget when  $\mathbb{E}[N] \leq \eta$ . It is said to be *deterministic* when every  $\pi_n$  takes values only in  $\{0, 1\}$  and it is said to be *randomized* otherwise. A randomized stopping rule can be seen as stochastic multiplexing among some number of deterministic stopping rules.

Our goal is to minimize the MSE in estimation of  $p$  through the design of a stopping rule that satisfies the trial budget and of an estimator  $\hat{p}(X_1, X_2, \dots, X_N)$ . We will first show that the continuation probability functions can be simplified greatly with no loss of optimality. Then, we will provide results on optimizing the stopping rule under a Beta prior on  $p$ .

#### A. Framework for Data-Dependent Stopping

Based on (9), a natural representation of a stopping rule is a node-labeled binary tree representing all sample paths of the Bernoulli process, with a probability of continuation label at each node. This representation has  $2^{d+1} - 1$  labels for observation sequences up to length  $d$ . However, the tree can be simplified to a trellis without loss of optimality. Conditioned on observing  $k$  successes in  $m$  trials, all  $\binom{m}{k}$  sequences of

<sup>3</sup>The time  $N$  does not satisfy the standard definition of a stopping time when the stopping rule is randomized because randomness independent of the sequence of outcomes  $\{X_n\}$  is allowed to influence the decision of whether or not to continue observations.




Fig. 4. A node-labeled trellis showing continuation probabilities for observation sequences up to length 5;  $q_{k,m}$  denotes the probability of continuing after observing  $k$  successes in  $m$  trials.

length  $m$  with  $k$  successes are equally likely. Thus, no improvement can come from having unequal continuation probabilities for the  $\binom{m}{k}$  tree nodes that each represent having  $k$  successes in  $m$  trials. Instead, these nodes should be combined, therefore reducing the tree to a trellis. This representation has  $\frac{1}{2}(d+1)(d+2)$  labels for observation sequences up to length  $d$ . The continuation probability functions are reduced to a set of probabilities  $\{q_{k,m} : m = 0, 1, \dots; k = 0, 1, \dots, m\}$  for continuing after  $k$  successes in  $m$  trials, as depicted in Fig. 4.

Trellises alone – without labels – give a simple representation for both data-dependent and data-independent deterministic stopping rules: Hence, we begin with some related terminology that will be used throughout this paper.

*Definition 1 (Complete trellis):* A complete trellis of depth  $d \in \mathbb{N}$  contains all nodes  $v = (k, m)$  belonging to the set

$$\mathcal{T}_d = \{(k, m) : k = 0, 1, \dots, m; m = 0, 1, \dots, d\}. \quad (10)$$

*Definition 2 (Strategy):* Any  $T \in 2^{\mathcal{T}_d}$  is a strategy when all nodes in  $T$  are connected and  $T$  contains the root node  $(0, 0)$ .

Henceforth, we restrict our attention to strategies and stochastic multiplexing among strategies. The stopping rule prescribed by the strategy  $T$  is

$$q_{k,m}(T) = \begin{cases} 1, & v = (k, m) \in T; \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

### B. Standard Sampling Methods and their Representations

The conventional use of a fixed number of trials  $n$  corresponds to continuation probabilities

$$q_{k,m} = \begin{cases} 1, & m < n; \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

Regardless of the sample path, one observes exactly  $n$  trials, and the number of successes  $K$  is a  $\text{Binomial}(n, p)$  random variable. We refer to this as *binomial sampling* or the *binomial stopping rule*. An example of the corresponding trellis representation for a fixed number of trials  $n = 5$  is shown in Fig. 5(a).

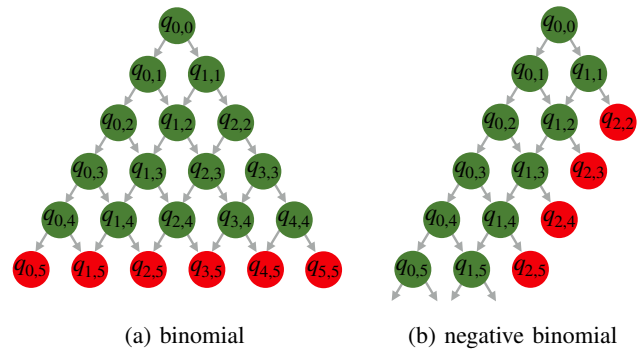


Fig. 5. Green nodes form the trellis representations  $T$  of (a) the binomial stopping rule with  $n = 5$  and (b) the negative binomial stopping rule with  $\ell = 2$ ; these are nodes with continuation probability 1. Red nodes are not in  $T$ ; these are nodes with continuation probability 0.

The technique analyzed by Haldane [7] and employed in first-photon imaging [14] with  $\ell = 1$  can be expressed with continuation probabilities

$$q_{k,m} = \begin{cases} 1, & k < \ell; \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

Observations cease with  $\ell$  successes in  $M$  trials, where  $M$  is a  $\text{NegativeBinomial}(\ell, p)$  random variable. We call such a strategy the *negative binomial stopping rule*, or *geometric stopping rule* for the special case where  $\ell = 1$ . The trellis representation of the negative binomial stopping rule for  $\ell = 2$  is shown in Fig. 5(b).

In general, observations cease with  $K$  successes in  $M$  trials, where  $K$  and  $M$  are both random variables. Importantly, the i.i.d. nature of a Bernoulli process makes the pair  $(K, M)$  contain all the information that is relevant from the sequence of observations. As noted in the reduction from tree to trellis, conditioned on  $(K, M) = (k, m)$ , all sequences of length  $m$  with  $k$  successes are equally likely, so the specific sequence among these is uninformative about  $p$ .

### C. Analysis Under Beta Prior

Our method for optimizing the design of continuation probabilities is through analyzing mean Bayes risk reduction from continuation. We define risk function  $L$  as *squared error* or *squared loss*

$$L(p, \hat{p}) = (p - \hat{p})^2,$$

where  $p$  is the unknown Bernoulli parameter and  $\hat{p}$  is the estimate of this parameter. The Bayes risk  $R$  is defined as

$$R(\hat{p}) = \mathbb{E}[L(p, \hat{p})] = \mathbb{E}[(p - \hat{p})^2],$$

which in this case is the MSE. Using the minimum MSE (MMSE) estimator, for which  $\hat{p} = \mathbb{E}[P]$ , the Bayes risk is the *variance* of the posterior distribution. Thus, key to the optimization is to track posterior variances through the trellis. For any prior on  $p$ , the posterior variance could be computed online or precomputed for some fixed trellis. Here we provide detailed computations only for the convenient case of choosing a conjugate prior.

1) *Beta Prior*: The Beta distribution is the conjugate prior for Bernoulli, binomial, and negative binomial distributions: When  $P$  has the Beta( $a, b$ ) distribution with probability density function

$$f_P(p; a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1},$$

where  $\Gamma(\cdot) = (\cdot - 1)!$  is the gamma function, the posterior distribution after observing  $k$  successes in  $m$  trials has the Beta( $a+k, b+m-k$ ) distribution. The beta distribution  $P \sim \text{Beta}(a, b)$  has mean

$$\mu_{a,b} = \mathbb{E}[P] = \frac{a}{a+b} \quad (14)$$

and variance

$$\sigma_{a,b}^2 = \text{var}(P) = \frac{ab}{(a+b)^2(a+b+1)}. \quad (15)$$

2) *Expected Number of Trials*: For the stopping rule represented by the trellis  $T$ , the expected number of trials is the weighted sum of the depths of all stopping (or leaf) nodes with weights corresponding to probability of reaching that node, under the initial prior. For a trellis  $T \in 2^{\mathcal{T}_d}$  and initial prior Beta( $\alpha, \beta$ ), the probability of reaching any node  $v = (k, m) \in \mathcal{T}_d$  can be expressed recursively using the probabilities of reaching its parents,  $(k-1, m-1)$  and  $(k, m-1)$ . Conditioned on reaching  $(k-1, m-1)$ , the probability of reaching  $(k, m)$  is the product of continuation probability  $q_{k-1, m-1}(T)$  and success probability

$$\mu_{\alpha+k-1, \beta+m-k} = \frac{\alpha+k-1}{\alpha+\beta+m-1}; \quad (16a)$$

similarly, conditioned on reaching  $(k, m-1)$ , the probability of reaching  $(k, m)$  is the product of continuation probability  $q_{k, m-1}(T)$  and failure probability

$$1 - \mu_{\alpha+k, \beta+m-k-1} = \frac{\beta+m-k-1}{\alpha+\beta+m-1}. \quad (16b)$$

Hence, we have the recursion

$$u_{k,m}(T) = u_{k-1, m-1}(T)q_{k-1, m-1}(T)\frac{\alpha+k-1}{\alpha+\beta+m-1} + u_{k, m-1}(T)q_{k, m-1}(T)\frac{\beta+m-k-1}{\alpha+\beta+m-1} \quad (17)$$

for the probability  $u_{k,m}(T)$  of reaching node  $(k, m)$ . The recursion is initialized with  $u_{0,0}(T) = 1$  and  $u_{k,m}(T) = 0$  when  $k \notin \{0, 1, \dots, m\}$ . Since  $T \in 2^{\mathcal{T}_d}$ , it suffices to compute up to  $m = d+1$ .

Using  $u_{k,m}(T)$  from (17), the expected number of trials incurred by a strategy  $T \in 2^{\mathcal{T}_d}$  starting with a Beta( $\alpha, \beta$ ) prior is

$$h_{\alpha, \beta}(T) = \sum_{v \in \mathcal{T}_{d+1} \setminus T} m u_{k,m}(T). \quad (18)$$

The nonzero terms in the sum correspond to the reachable leaf nodes, which are all contained in  $T' = \mathcal{T}_{d+1} \setminus T$ .

3) *Expected Bayes Risk*: Under initial prior Beta( $\alpha, \beta$ ), the Bayes risk of the estimate of  $p$  from observations leading to node  $(k, m)$  is given by (15), with  $a = \alpha + k$  and  $b = \beta + m - k$ . A strategy  $T$  has expected Bayes risk  $g_{\alpha, \beta}(T)$  given by the sum of the Bayes risks of nodes with zero continuation probability weighted by the probabilities of reaching that node:

$$g_{\alpha, \beta}(T) = \sum_{v \in T'} u_{k,m}(T) \sigma_{\alpha+k, \beta+m-k}^2 = \sum_{v \in T'} u_{k,m}(T) \frac{(\alpha+k)(\beta+m-k)}{(\alpha+\beta+m)^2(\alpha+\beta+m+1)}. \quad (19)$$

4) *Optimization Problem Statement*: With the proposed trellis-based framework, finding an optimal deterministic stopping rule (in the MSE sense) under an average budget constraint becomes a set minimization problem:

$$T^* = \arg \min_{T \in 2^{\mathcal{T}_d}} g_{\alpha, \beta}(T) \quad (20)$$

subject to  $h_{\alpha, \beta}(T) \leq \eta$ .

Implementable solutions to (20), with varying complexities and deviations from optimality, are presented in the subsequent section. We seek only solutions on the lower convex hull of the trade-off between  $\eta$  and  $\min g_{\alpha, \beta}(T)$ . Stochastic multiplexing among these solutions gives optimal randomized stopping rules.

## IV. STOPPING RULE DESIGN

### A. A Dynamic Programming Solution

For a fixed and sufficiently large  $d$ , total enumeration of the entire solution space is a possible approach for solving (20) to find an optimal deterministic stopping rule. However, the combinatorial structure of the problem means that evaluating the Bayes risks (19) and expected numbers of trials (18) for all possible strategies can be computationally prohibitive, even for moderate trial budgets; this precludes full enumeration.

Conversely, one could start at the leaf nodes of a complete trellis (with depth  $d$ ), traverse the trellis towards its root, whilst deciding whether each visited node merits inclusion in the optimized solution. This is the basis of a dynamic programming (DP) solution: it solves our optimization problem that involves making a sequence of decisions by determining, for each decision, subproblems that can be solved in a similar fashion [33]. As such, a solution of the original problem can be found from solutions of subproblems.

Precisely, we first relax (20) by writing its Lagrangian formulation:

$$\min_{T \in 2^{\mathcal{T}_d}} g_{\alpha, \beta}(T) + \lambda h_{\alpha, \beta}(T), \quad (21)$$

where  $\lambda \in \mathbb{R}_+$ , can be viewed as the desired MSE reduction per additional trial. We introduce three compact notations associated with node  $v = (k, m)$ :

$$R_{k,m}^{\alpha, \beta} = \sigma_{\alpha+k, \beta+m-k}^2 = \frac{(\alpha+k)(\beta+m-k)}{(\alpha+\beta+m)^2(\alpha+\beta+m+1)} \quad (22a)$$

is the mean Bayes risk conditioned on stopping at  $v$ ,

$$S_{k,m}^{\alpha, \beta} = \mu_{\alpha+k, \beta+m-k} = \frac{\alpha+k}{\alpha+\beta+m} \quad (22b)$$



is the probability of the next trial being a success, and

$$F_{k,m}^{\alpha,\beta} = 1 - \mu_{\alpha+k,\beta+m-k} = \frac{\beta + m - k}{\alpha + \beta + m} \quad (22c)$$

is the probability of the next trial being a failure. The dynamic program summarized in Algorithm 1 iteratively constructs a solution to (21) by comparing  $R_{k,m}^{\alpha,\beta}$  to the lowest cost achievable from the state that results after a single trial. More precisely, at each node  $(k, m)$ ,  $R_{k,m}^{\alpha,\beta}$  is compared to the cost of one additional trial plus the expected lowest cost achievable from the subsequent state. We keep track of the lowest of these values in  $V_{k,m}$ , which is the lowest achievable cost from any node  $(k, m)$ . If  $R_{k,m}^{\alpha,\beta}$  is lower, then an additional trial is not warranted and the node is eliminated, i.e.  $T \leftarrow T \setminus \{v\}$  and  $q_{k,m} = 0$ . Because of the decomposability of the problem, the solutions are optimal.

---

**Algorithm 1** Dynamic programming algorithm to find optimal deterministic stopping rule for Beta( $\alpha, \beta$ ) prior, Lagrange multiplier  $\lambda$ , and maximum depth  $d$

---

**Input:**  $(\alpha, \beta)$ ,  $\lambda \in \mathbb{R}_+$ ,  $d \in \mathbb{N}$   
**Output:**  $T^* \in 2^{\mathcal{T}_d}$   
**Initialize:**  $T = \mathcal{T}_d$  and  $[V_{k,m}]_{k,m} = 0$  for all  $k, m$   
**for**  $k = 1, \dots, d$  **do**  
    Set  $V_{k,d} \leftarrow R_{k,d}^{\alpha,\beta}$  using (22a)  
**end for**  
**for**  $m = d - 1, \dots, 1$  **do**  
    **for**  $k = 1, 2, \dots, m$  **do**  
        **if**  $R_{k,m}^{\alpha,\beta} > \lambda + S_{k,m}^{\alpha,\beta} V_{k+1,m+1} + F_{k,m}^{\alpha,\beta} V_{k,m+1}$  **then**  
             $V_{k,m} \leftarrow \lambda + S_{k,m}^{\alpha,\beta} V_{k+1,m+1} + F_{k,m}^{\alpha,\beta} V_{k,m+1}$   
        **else**  
             $V_{k,m} \leftarrow R_{k,m}^{\alpha,\beta}$   
             $T \leftarrow T \setminus \{v\}$   
        **end if**  
    **end for**  
**end for**  
**return**  $T^* \leftarrow T$

---

### B. A Greedy Algorithm

The DP method (Algorithm 1) prunes from the complete trellis  $\mathcal{T}_d$ . Monotonicity of the objective  $g_{\alpha,\beta}(T)$  and cost  $h_{\alpha,\beta}(T)$  can be exploited to develop a lower-complexity greedy algorithm that instead builds a trellis starting from just the root node.

The scheme outlined in Algorithm 2 monotonically improves the objective function value for the minimization problem (20) with each iteration. Specifically, at iteration  $i$ , the greedy decision is to add to the current trellis  $T_i$  a node  $v \notin T_i$  that yields the largest reduction in the Bayes risk *per additional trial*,

$$\frac{g_{\alpha,\beta}(T_i \cup v) - g_{\alpha,\beta}(T_i)}{h_{\alpha,\beta}(T_i \cup v) - h_{\alpha,\beta}(T_i)},$$

without violating the mean number of trials constraint. The scheme terminates when no such node exists.

---

**Algorithm 2** Greedy algorithm to find deterministic stopping rule for Beta( $\alpha, \beta$ ) prior and trial budget  $\eta$

---

**Input:**  $(\alpha, \beta)$ ,  $\eta$   
**Output:**  $T^*$   
**Initialize:**  $i \leftarrow 0$ ,  $T_0 \leftarrow \{\}$   
**repeat**  
     $\tilde{v} \leftarrow \arg \min_{v \notin T_i} \frac{g_{\alpha,\beta}(T_i \cup v) - g_{\alpha,\beta}(T_i)}{h_{\alpha,\beta}(T_i \cup v) - h_{\alpha,\beta}(T_i)}$   
     $T_{i+1} \leftarrow T_i \cup \tilde{v}$   
     $i \leftarrow i + 1$   
**until**  $h_{\alpha,\beta}(T_i) > \eta$   
**return**  $T^* \leftarrow T_{i-1}$

---

### C. Online Threshold-Based Termination

Our final method applies a simple rule for termination of trials, depending on the prior parameters  $(\alpha, \beta)$  and the  $(k, m)$  position in the trellis. It implies a trellis design, but it does not require storage of a designed trellis.

Suppose a sequence of trials reaches a node corresponding to the posterior distribution Beta( $a, b$ ). Denote the mean Bayes risk *without* performing an additional trial by

$$R_{\text{stop}}(a, b) = \sigma_{a,b}^2, \quad (23)$$

using the variance given in (15). When one additional trial is performed, the posterior distribution is either Beta( $a + 1, b$ ) if the outcome of the additional trial is a success, or Beta( $a, b + 1$ ) if the outcome of the additional trial is a failure. Therefore, the mean Bayes risk resulting from continuing with one additional trial is

$$\begin{aligned} R_{\text{cont}}(a, b) &= \mathbb{E}[(1 - P) \sigma_{a,b+1}^2 + P \sigma_{a+1,b}^2] \\ &= \frac{ab}{(a + b)(a + b + 1)^2}. \end{aligned} \quad (24)$$

The Bayes risk *reduction* from one additional trial is

$$\begin{aligned} \Delta R(a, b) &= R_{\text{stop}}(a, b) - R_{\text{cont}}(a, b) \\ &= \frac{ab}{(a + b)^2(a + b + 1)^2}. \end{aligned} \quad (25)$$

Recall that, starting from a Beta( $\alpha, \beta$ ) prior, upon reaching node  $(k, m)$ , the posterior is Beta( $\alpha + k, \beta + m - k$ ). The Bayes risk reduction from an additional trial,

$$\Delta R(k, m; \alpha, \beta) = \frac{(\alpha + k)(\beta + m - k)}{(\alpha + \beta + m)^2(\alpha + \beta + m + 1)^2}, \quad (26)$$

can be the basis of an online stopping rule. Let  $\Delta_{\min} > 0$  denote a specified threshold value for the reduction in Bayes risk that justifies an additional trial. Then stopping based on this threshold induces the probabilities of continuing at each node of the trellis given by

$$q_{k,m} = \begin{cases} 1, & \Delta R(k, m; \alpha, \beta) \geq \Delta_{\min}; \\ 0, & \Delta R(k, m; \alpha, \beta) < \Delta_{\min}. \end{cases} \quad (27)$$

Fig. 6(a) shows values of  $\Delta R(k, m; 1, 1)$  for  $m = 0, 1, \dots, 5$ . The choice of threshold  $\Delta_{\min} = 0.005$  results in the trellis shown in Fig. 6(b).

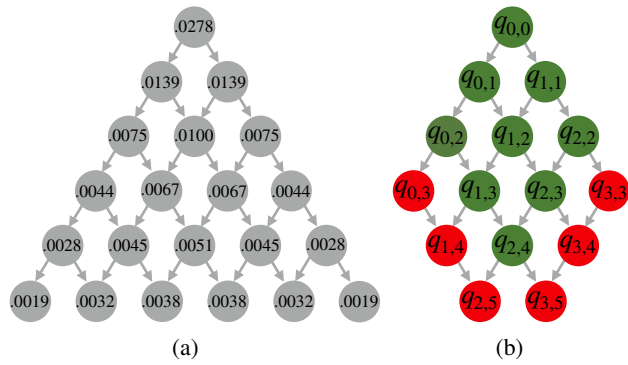


Fig. 6. (a) Bayes risk reductions per additional trial. (b) Resulting trellis of continuation probabilities for  $\Delta_{\min} = 0.005$  (right). Beta(1, 1) (i.e., uniform) prior for  $P$  has been assumed.

Notice that for a fixed trellis depth  $m$ , the denominator of (26) is fixed, and the numerator of (26) is a product of factors with fixed sum that is equal to  $\alpha + \beta + m$ . Thus, from the arithmetic–geometric mean inequality,  $\Delta R(k, m; \alpha, \beta)$  is largest where the posterior distribution is most symmetric. This is apparent in the example in Fig. 6(b); since we have started with a uniform prior, the center of each row represents a symmetric posterior, and additional observations are most merited near the center of each row. Starting with a highly asymmetric prior ( $\alpha \ll \beta$  or  $\alpha \gg \beta$ ), the same principle explains an asymmetry in the greedily optimized trellis of continuation probabilities.

*Example 4 (Suboptimality of binomial sampling):* Suppose we have a Beta(1, 1) (i.e., uniform) prior. Then (26) simplifies to

$$\Delta R(k, m; 1, 1) = \frac{(k+1)(m-k+1)}{(m+2)^2(m+3)^2}.$$

For the threshold-based termination to induce binomial sampling with  $m^*$  trials, the incremental benefit  $\Delta R$  at  $(k, m) = (0, m^*)$  must be greater than  $\Delta R$  at  $(k, m) = (\lfloor \frac{1}{2}(m^* + 1) \rfloor, m^* + 1)$ :

$$\frac{m^* + 1}{(m^* + 2)^2(m^* + 3)^2} \geq \frac{(\lfloor \frac{1}{2}(m^* + 1) \rfloor + 1)(m^* - \lfloor \frac{1}{2}(m^* + 1) \rfloor + 2)}{(m^* + 3)^2(m^* + 4)^2}. \quad (28)$$

Since (28) fails to hold for any  $m^* > 2$ , threshold-based termination induces binomial sampling only for 1 and 2 trials. This is consistent with Fig. 6. For such a small trial budget, full enumeration of stopping rules is also feasible, and one can conclude that binomial sampling is indeed suboptimal for any trial budget greater than 2. Similar arguments can be made for non-uniform beta priors.

#### D. Comparisons of Designs

Sweeping  $\Delta_{\min}$  in threshold-based termination is very similar to sweeping  $\lambda$  in Algorithm 1; it will achieve certain mean numbers of trials, similar to sweeping  $\eta$  in Algorithm 2. Intermediate values of the mean number of trials can be achieved by finding  $(k^*, m^*)$  such that  $\Delta R(k, m; \alpha, \beta)$  is largest among those below  $\Delta_{\min}$  and varying  $q_{k^*, m^*}$  over




Fig. 7. Dynamic programming solution minus online threshold-based termination result, both with mean number of trials  $\approx 95.36$ . Online threshold-based termination achieves mean Bayes risk of 0.0016037 whereas DP gives 0.0016036. At 24 nodes (red, ‘-1’), threshold-based rule performs additional trials and DP does not; at 54 nodes (blue, ‘+1’), DP performs trials and threshold-based rule does not. The greedily designed trellis coincides with the DP trellis, hence their difference plot is omitted.




Fig. 8. Comparison between proposed strategies. DP and greedy algorithm solutions coincide for all trial budgets, while threshold-based termination is only very slightly worse by a factor of at most 1.000195.

$(0, 1)$ . This idea is used to enforce an equal expected number of trials for trellises optimized with each method, thus allowing a fair comparison of their Bayes risks. For a mean number of trials  $\approx 95.36$ , Algorithms 1 and 2 were found to give exactly the same trellis, while online threshold-based termination gave a slightly different trellis with slightly higher mean Bayes risk. Fig. 7 illustrates the difference in  $q_{k,m}$  values. It is zero for the vast majority of nodes, with 24 nodes at which the DP-designed trellis terminates but the threshold-based rule does not (red,  $-1$ ), and 54 nodes at which the threshold-based rule terminates but the DP-designed trellis does not (blue,  $+1$ ).<sup>4</sup>

Illustrated in Fig. 8 is a comparison of our three proposed implementable strategies, applied for a uniform prior, over a range of trial budgets. MSEs of DP (Algorithm 1) and the greedily optimized trellis (Algorithm 2) coincide for all trial budgets because the trellises are identical – though we have not proven that this is guaranteed. The online threshold-based stopping rule is only very slightly worse by a factor of at most 1.000195 (less than 0.001 dB).

The phenomenon of more trials being merited when  $p$  is near  $\frac{1}{2}$  counteracts the MSE of  $p(1-p)/n$  being largest for  $p$  near  $\frac{1}{2}$ . This is illustrated in Fig. 9(a), which shows mean numbers of trials allocated as a function of  $p$ . We have optimized for MSE averaged over  $p$  and, in so doing,

<sup>4</sup>The mean number of trials is equal. To be convinced that the blue and red nodes can balance, note that while there are more blue nodes, they are for larger values of  $m$  and thus have lower probabilities of being reached.



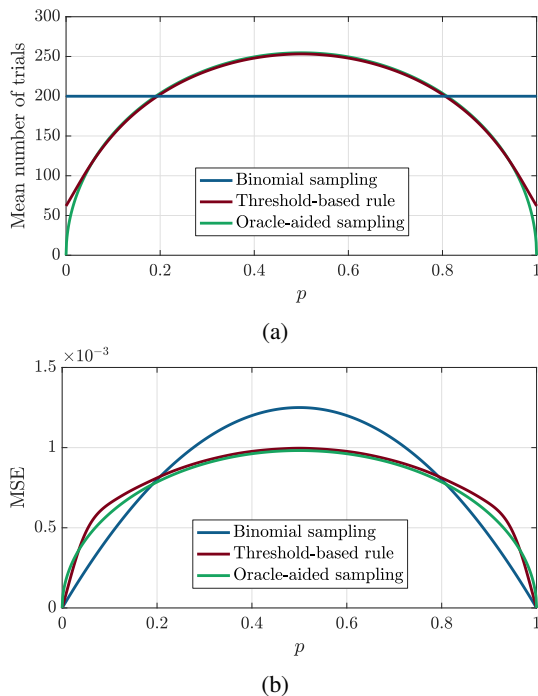


Fig. 9. Dependences on  $p$  under a uniform prior. (a) Conditional expectation of the number of trials, conditioned on  $p$ , for binomial sampling, online threshold-based termination, and the oracle-aided binomial stopping rule. Expected number of trials is 200 for all of the three methods. (b) Dependence of conditional MSEs on the true Bernoulli parameter  $p$  under a trial budget of  $\eta = 200$ .

obtained a modest improvement factor of  $\approx 1.05$  in this average, comparing the online threshold-based termination to conventional binomial sampling. A more significant reduction in the worst-case MSE is a by-product of the optimization (see Fig. 9(b)).

### E. Asymptotic Comparison with Oracle-Aided Allocation

Considering the non-degenerate cases  $p \in (0, 1)$ , the threshold-based termination is asymptotically equivalent to oracle-aided optimal allocation. For a large trial budget  $\eta$ , we will find an approximation for  $m^\circ$ , the number of trials at which the online threshold-based rule terminates. This will match the form of (2) or (6).

Using (26), for an initial Beta( $\alpha, \beta$ ) prior, the online rule continues at node  $(k, m)$  if and only if

$$\frac{(\alpha + k)(\beta + m - k)}{(\alpha + \beta + m)^2(\alpha + \beta + m + 1)^2} \leq \Delta_{\min}. \quad (29)$$

Since the trial budget is large and  $p \in (0, 1)$ ,  $k$ ,  $m$ , and  $m - k$  are all large when nearing termination. Hence, we approximate

the expression in (29) as

$$\begin{aligned} & \frac{(\alpha + k)(\beta + m - k)}{(\alpha + \beta + m)^2(\alpha + \beta + m + 1)^2} \\ &= \frac{m^2(\alpha/m + k/m)(\beta/m + 1 - k/m)}{m^4(\alpha/m + \beta/m + 1)^2(\alpha/m + \beta/m + 1 + 1/m)^2} \\ &\approx \frac{(k/m)(1 - k/m)}{m^2} \\ &= \frac{\hat{p}_{\text{ML}}(1 - \hat{p}_{\text{ML}})}{m^2}, \end{aligned} \quad (30)$$

where  $\hat{p}_{\text{ML}} = k/m$  is the ML estimate of  $p$ .

Substituting (30) into (29), we obtain

$$m^\circ \approx \sqrt{\frac{\hat{p}_{\text{ML}}(1 - \hat{p}_{\text{ML}})}{\Delta_{\min}}}. \quad (31)$$

By the law of large numbers,  $\hat{p}_{\text{ML}} \rightarrow p$ , so (31) shows a match to (2), with  $\Delta_{\min}$  determining the trial budget. Furthermore, by comparison with (6), we see an equivalence by choosing  $\Delta_{\min} = \mathbb{E}[\sqrt{P(1-P)}]/\eta$ .

Fig. 9(a) illustrates an example of the approximate match between threshold-based termination and oracle-aided sampling that is predicted by the match among (2), (6), and (31). Note that convergence is not uniform in  $p$ ; a larger trial budget is needed to observe approximate equivalence in allocations for  $p$  near 0 and near 1.

Fig. 10 shows the variation of the MSEs with mean number of trials budget constraint for conventional binomial sampling, threshold-based termination, and oracle-aided allocation. The results are based on Monte Carlo simulations, with MATLAB, using the phantom image in Fig. 3(a). As expected the optimized rules consistently achieve MSE improvements over the conventional binomial sampling, for all simulated trial budgets. In addition, when compared to the unrealizable oracle-aided method, the threshold-based approach only marginally under-performs at moderate mean number of trials budget constraints. This observation is further underscored in Fig. 9, which show significant overlap between threshold-based termination and oracle-aided allocation, in terms of both trial allocations and the resulting MSEs. Using a negative binomial sampling strategy yields significantly worse performance than binomial sampling and our proposed rules for estimating  $p$ ; thus, we have omitted it from Figs. 9 and 10, as well as other numerical simulations related to the estimation of  $p$ .

In line with the earlier asymptotic analysis, the threshold-based termination and oracle-aided performances coincide for moderate to high mean numbers of trials, independent of the prior. Under a highly skewed prior consistent with the true distribution of the phantom pixels, Fig. 10(b) demonstrates that it is possible for online threshold-based termination to even outperform the oracle-aided binomial method, at low trial budgets. This phenomenon is attributable to the online method allocating more trials when the Bernoulli process realization has a relatively high fraction of successes. Put simply, it is allocating more trials for “unlucky” realizations where the MSE would be higher, while the oracle-aided binomial method maintains a fixed number of trials.

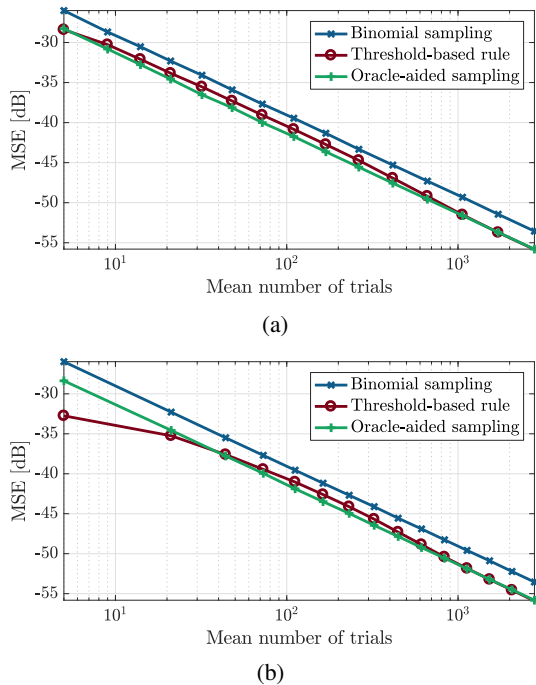


Fig. 10. Results for the Shepp–Logan Phantom scaled to [0.001, 0.101], assuming: (a) Beta(1, 1) and (b) Beta(1, 50), for the online threshold-based termination. The MSE has been computed from the average of 20 independent experiments for each mean number of trials. Improvements are consistent with the trial allocation gain computed in Example 2(d).

## V. ESTIMATING FUNCTIONS OF A BERNOULLI PARAMETER

When estimating an arbitrary function  $f(p)$  of a Bernoulli parameter is of interest [11], one can derive similar stopping strategies as before. In this section, we concern ourselves only with the estimation of  $f(p) = \log p$  due to its prevalence in real-life scenarios. For instance, the subjective brightness perceived by the human vision system is a logarithmic function of the incident light intensity [34], [35]. Also, the common *log odds ratio*  $\log(p/(1-p))$ , is approximately equal to  $\log p$  when  $p \ll 1$ .

As before, we begin with a squared error loss

$$L(p, \cdot) = (f(p) - \hat{f}(\cdot))^2, \quad (32)$$

where  $\hat{f}(\cdot)$  is the estimate of  $f(p)$ . The expectation of this loss function over  $p$  gives the Bayes risk

$$R(\cdot) = \mathbb{E}[L(p, \cdot)] = \mathbb{E}[(f(p) - \hat{f}(\cdot))^2]. \quad (33)$$

For  $f(p) = \log p$ , suppose a sequence of trials leads to a node in the trellis corresponding to the posterior distribution Beta( $a, b$ ). Under  $P \sim \text{Beta}(a, b)$ , the MMSE estimator of  $\log p$  is [36]

$$\hat{f}(a, b) := \mathbb{E}[\log P] = \psi^{(0)}(a) - \psi^{(0)}(a + b), \quad (34)$$

where  $\psi^{(m)}$  is the polygamma function of order  $m$ . The Bayes risk (33) when no additional trial is performed,  $R_{\text{stop}}$ , becomes the variance of  $\log p$  [37]:

$$R_{\text{stop}} \stackrel{\text{def}}{=} R(a, b) = \text{var}(\log P) = \psi^{(1)}(a) - \psi^{(1)}(a + b). \quad (35)$$




Fig. 11. Bayes risk reductions for estimation of (a)  $f(p) = p$  and (b)  $f(p) = \log p$ , assuming Beta(1, 1) assumed initially. Low  $p$  values are assigned significantly more trials when  $\log p$  is estimated.

If one additional trial is performed, the Bayes risk reduces to

$$\begin{aligned} R_{\text{cont}}(a, b) &= \mathbb{E}[(1 - P) R(a, b + 1) + P R(a + 1, b)] \\ &= \frac{b}{a + b} R(a, b + 1) + \frac{a}{a + b} R(a + 1, b). \end{aligned} \quad (36)$$

Hence, the Bayes risk reduction from one additional trial is

$$R_{\text{stop}}(a, b) - R_{\text{cont}}(a, b) = \frac{b}{a(a + b)^2}. \quad (37)$$

Starting with prior Beta( $\alpha, \beta$ ), the counterpart to (26) for estimation of  $\log p$  is

$$\Delta R(k, m; \alpha, \beta) = \frac{\beta + m - k}{(\alpha + k)(\alpha + \beta + m)^2}. \quad (38)$$

As before, this can be used in (27) as an online threshold-based termination method.

The Bayes risk reductions for both  $f(p) = p$  in (26) and  $f(p) = \log p$  in (38), starting with a uniform prior, are shown as heat maps in Fig. 11. When  $f(p) = \log p$ , the reduction from additional trials after observing sequences with low number of successes is significantly larger. Thus, the online threshold-based termination of Section IV-C is likely to assign more trials for the smaller underlying Bernoulli parameters. Such a stopping rule is intuitive because a fixed amount of estimation error for  $p$  would contribute more to the loss function defined in (32), with  $f(p) = \log p$ , when  $p$  is small. In fact, one can choose a loss function which enforces different penalties for different  $p$  values. An example is the family of weighted mean squared errors,  $\mathbb{E}[w(p)(p - \hat{p})^2]$ , where a weighting function  $w(p)$  is designed according to the problem. A special case is relative MSE  $\mathbb{E}[(p - \hat{p})^2/p^2]$ , which is approximately the squared loss in (33) with  $f(p) = \log p$  for estimates sufficiently close to the true value [11].

When smaller  $p$  values are of more importance, it makes sense to use a strategy that allocates more trials to these instances. Negative binomial sampling explained in Section III-B




Fig. 12. Estimation of  $\log p$ . Conditional MSE of  $\log p$ , conditioned on  $p$ , is shown as a function of  $p$  for binomial sampling ( $\eta = 72$ ), negative binomial sampling ( $\ell = 10$ , inducing  $\eta = 71.66$ ), and online threshold-based termination ( $\eta = 71.77$ ). The threshold-based termination uses a uniform prior, and averaging over a uniform prior gives MSE of 0.0668 for binomial sampling, 0.0489 for negative binomial sampling, and 0.0328 for threshold-based termination. These values are computed directly from their corresponding trellises (hence exact), not through numerical integration.

achieves this type of trial allocation. For the estimation of  $f(p) = \log p$ , we compare the performances of binomial sampling, negative binomial sampling, and online threshold-based termination in Fig. 12. Threshold-based termination outperforms both binomial and negative binomial sampling, with improvement factors of 2.037 and 1.491, respectively.

## VI. APPLICATIONS TO ACTIVE IMAGING

Active imaging systems typically raster scan the scene by probing patch  $(i, j)$ ,  $i = 1, \dots, N_i$  and  $j = 1, \dots, N_j$ , using pulsed illumination. The measured data – used to form an image of the scene – are arrays  $[k_{i,j}]_{i,j}$  and  $[m_{i,j}]_{i,j}$ ; i.e., the number of detections (successes) and number of illumination pulses (trials) for each scene patch. Note that the conventional approach of a fixed number of trials makes  $m_{i,j} = \eta$  for all  $(i, j)$  and  $\{k_{i,j}\}$  random, whereas both  $\{k_{i,j}\}$  and  $\{m_{i,j}\}$  are random when the proposed approach is applied.

The parameters of the Bernoulli processes generated by probing a scene patch and its neighbors are typically correlated. This can be exploited in the image formation stage through mechanisms inspired by any of various image compression or denoising methods. For this initial demonstration of adaptive acquisition, we apply total variation (TV) regularization [38]. We present simulation results using the Shepp–Logan phantom in Fig. 3(a), two lidar datasets provided by the Alaska Department of Natural Resources [39], and scanning electron microscopy (SEM) images *Foraminifera*<sup>5</sup> and *HairStyle*<sup>6</sup> taken from ThermoFisher Scientific. All images have been rescaled to take on values in the range  $[0.001, 0.101]$ .

### A. Estimation of $f(p) = p$

We focus here on comparing conventional binomial sampling against online threshold-based termination (27) applied

<sup>5</sup>Quanta SEM image of Protozoan group secreting a calcareous shell by Philippe Crassous, <https://www.fei.com/image-gallery/Foraminifera-Protozoan/>

<sup>6</sup>Quanta SEM image of the upper part of the style and stigma from an *Arabidopsis* flower by Guichuan Hou, [https://www.fei.com/uploadedImages/FEISite/Content/Image\\_Gallery/Images/2013\\_Image\\_Constest/FEI/IM\\_20130718\\_Hou\\_18\\_HairStyle\\_lg.jpg](https://www.fei.com/uploadedImages/FEISite/Content/Image_Gallery/Images/2013_Image_Constest/FEI/IM_20130718_Hou_18_HairStyle_lg.jpg)

TABLE I  
AVERAGE RECONSTRUCTION PSNRs, AVERAGED OVER 100 EXPERIMENTS, FOR CONVENTIONAL BINOMIAL SAMPLING AND ONLINE THRESHOLD-BASED TERMINATION, FOR TRIAL BUDGETS  $\eta$ .

Image	$\eta$	Pixelwise MMSE estimation		TV+ML estimation	
		Binomial	Threshold-based	Binomial	Threshold-based
Shepp–Logan	100	17.5 dB	<b>20.0 dB</b>	21.9 dB	<b>26.3 dB</b>
	200	19.9 dB	<b>22.1 dB</b>	24.9 dB	<b>28.5 dB</b>
lidar #1	400	18.1 dB	<b>18.8 dB</b>	24.0 dB	<b>25.9 dB</b>
	800	21.4 dB	<b>22.3 dB</b>	25.2 dB	<b>27.7 dB</b>
lidar #2	400	16.9 dB	<b>18.3 dB</b>	26.2 dB	<b>26.8 dB</b>
	800	20.7 dB	<b>22.3 dB</b>	27.6 dB	<b>28.8 dB</b>
<i>Foraminifera</i>	400	19.2 dB	<b>21.2 dB</b>	25.5 dB	<b>26.7 dB</b>
	500	20.4 dB	<b>22.4 dB</b>	25.2 dB	<b>27.5 dB</b>
<i>HairStyle</i>	400	21.2 dB	<b>22.7 dB</b>	26.8 dB	<b>27.8 dB</b>
	500	22.3 dB	<b>23.8 dB</b>	27.8 dB	<b>28.5 dB</b>

for each pixel. For  $f(p) = p$ , the relevant Bayes risk reduction per trial is given by (26).

1) *MMSE Estimation Under I.I.D. Prior*: When not exploiting any spatial correlations, each pixel estimation is performed separately using the methods of Section IV-C. Under a  $\text{Beta}(\alpha, \beta)$  prior, the MMSE estimate is  $\hat{p}_{\text{MMSE}}[i, j] = (k_{i,j} + \alpha) / (m_{i,j} + \alpha + \beta)$ .

Fig. 13(a) shows that with the choice of the  $\text{Beta}(2, 152)$  prior, MSE improvement of 2.42 dB is attained for trial budget  $\eta = 200$  for the Shepp–Logan phantom. In Figs. 13(b)–(e), MSE improvements were also demonstrated for the lidar and SEM images under various trial budgets and initial prior parameters. Improvements ranging from 0.92 dB for lidar #1 to 2.02 dB for *Foraminifera* were obtained. For the same corresponding choices of prior, we also perform 100 independent experiments for each test image at each of two different trial budgets; the results indicated in Table I show similar improvement factors. In particular, the performance gains are close to the prediction from the trial allocation gain (Example 2(d) for the Shepp–Logan phantom image). The same is observed for other images. Furthermore, we show in Fig. 14 that significant MSE improvements can be attained for a large range of Beta priors, using the Shepp–Logan phantom and *HairStyle* datasets. We also observe that the performance of binomial sampling is degraded more by a mismatched prior than the performance of threshold-based termination.

2) *TV-Regularized ML Estimation*: Reconstruction quality can be improved through the use of TV-regularized ML estimation [1], [38]. In one typical experimental trial shown in Fig. 15(a), the TV-regularized reconstruction from data obtained with online threshold-based termination outperforms the conventional binomial sampling by 4.36 dB in MSE for the Shepp–Logan phantom; the trial budget of  $\eta = 200$  and prior of  $\text{Beta}(2, 152)$  are the same as used previously. As anticipated, Figs. 15(b)–(e) also demonstrate significant improvements in MSE, ranging from 1.15 dB to 4.17 dB, for the remaining test images. Furthermore, keeping the corresponding priors and trial budgets used for each test image in Figs. 15, Table I also provides results averaged over 100 experiments for statistical significance. In many cases, imposing TV regularization increases the performance gained from the data-adaptive stopping rule. Most importantly, it does not completely diminish the gain of adapting the acquisition.

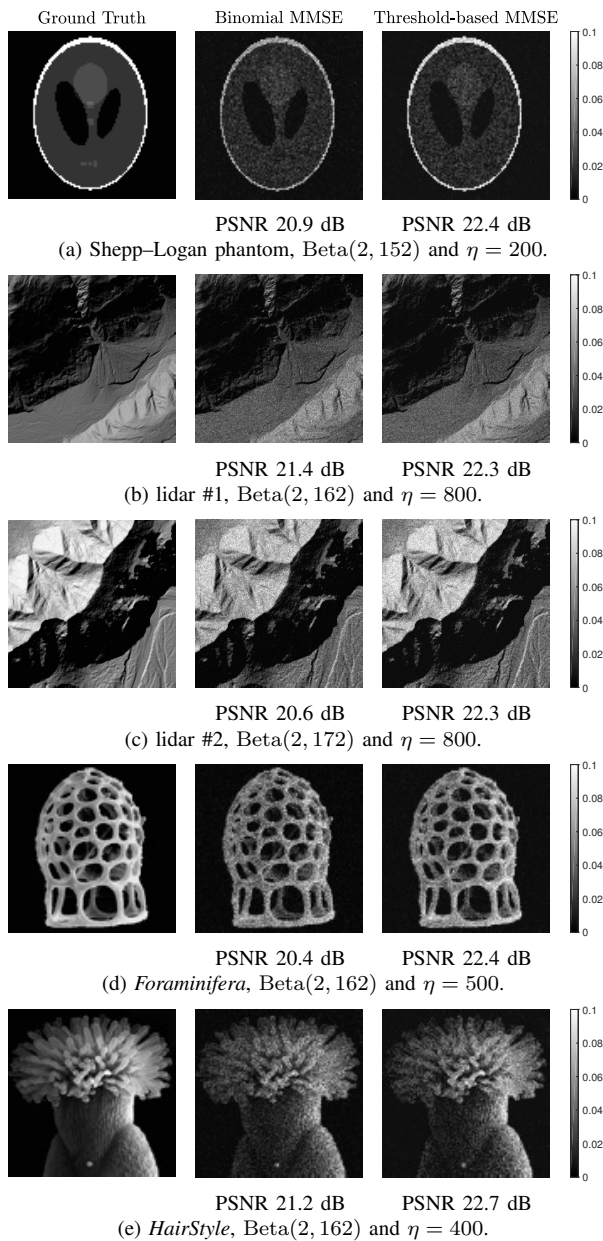


Fig. 13. Images reconstructed through pixelwise MMSE estimation showing MSE improvements of up to 2.42 dB for online threshold-based termination in place of conventional binomial sampling. Assumed priors and trial budgets are indicated below each test image. All images are scaled to  $[0.001, 0.101]$ .

### B. Estimation of $f(p) = \log p$

Now we present simulation results for the estimation of the logarithm of the previous test images. For  $f(p) = \log p$ , the Bayes risk reduction per trial to use in online threshold-based termination (27) is given by (38). Fig. 16 shows simulation results wherein improvement factors of 1.48 dB to 1.86 dB are observed using threshold-based termination compared with binomial sampling, and 2.56 dB to 3.78 dB when using threshold-based termination over negative binomial. Note that since the contributions to the error from the lower pixel values are higher, the comparison is provided at much higher trial budgets (i.e.,  $\eta = 3000, 1600, 1700, 2800,$  and  $1800$  for Figs. 16(a)–(e), respectively) than in Section VI-A to obtain

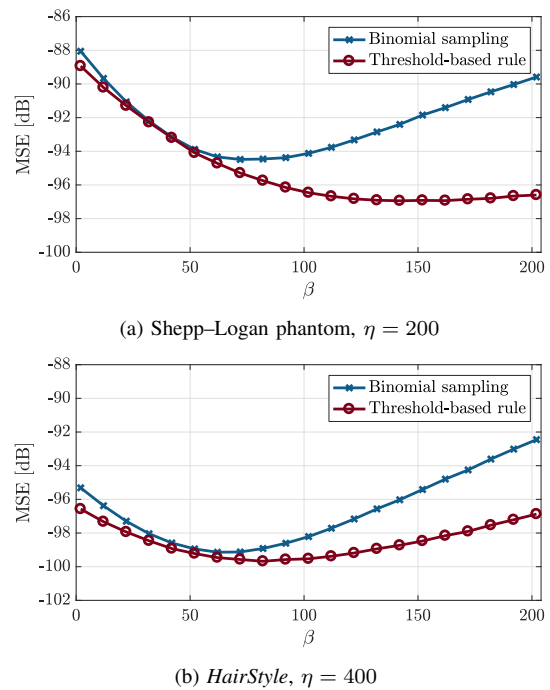


Fig. 14. Dependence of MSE on  $\beta$  when Beta( $2, \beta$ ) prior is assumed and pixelwise MMSE estimation is performed; each plotted point is obtained by averaging over 100 independent experiments.

TABLE II  
ESTIMATION OF  $f(p) = \log(p)$ . AVERAGE RECONSTRUCTION MSEs, AVERAGED OVER 100 EXPERIMENTS, FOR CONVENTIONAL BINOMIAL SAMPLING, NEGATIVE BINOMIAL SAMPLING AND THRESHOLD-BASED TERMINATION, FOR TRIAL BUDGETS  $\eta$ .

Image	$\eta$	Method		
		Binomial	Neg. Binomial	Threshold-based
Shepp-Logan	1800	0.271	0.481	<b>0.227</b>
	3000	0.208	0.253	<b>0.139</b>
lidar #1	1600	0.156	0.246	<b>0.113</b>
	2200	0.133	0.165	<b>0.080</b>
lidar #2	1700	0.154	0.237	<b>0.099</b>
	2300	0.135	0.161	<b>0.073</b>
Foraminifera	1700	0.253	0.462	<b>0.208</b>
	2700	0.206	0.245	<b>0.131</b>
HairStyle	1800	0.166	0.247	<b>0.118</b>
	2500	0.143	0.166	<b>0.087</b>

meaningful results. Dark regions having values of 0.001, for instance, require 1000 trials on average to observe a success. The effect of increasing number of trials is apparent in Table II. An increase in the improvement factor obtained by using the threshold-based rule, compared to both binomial and negative binomial stopping is observed, as the trial budget is increased. For example, for the Shepp-Logan phantom, an increase in the improvement factor from 0.77 dB to 1.75 dB is observed when increasing  $\eta$  from 1800 to 3000, if the threshold-based rule is used over conventional binomial stopping. An even larger increase is obtained when the threshold-based rule is compared against negative binomial stopping. The trend of increasing MSE improvements with increased trial budgets is persistent for other test images too.



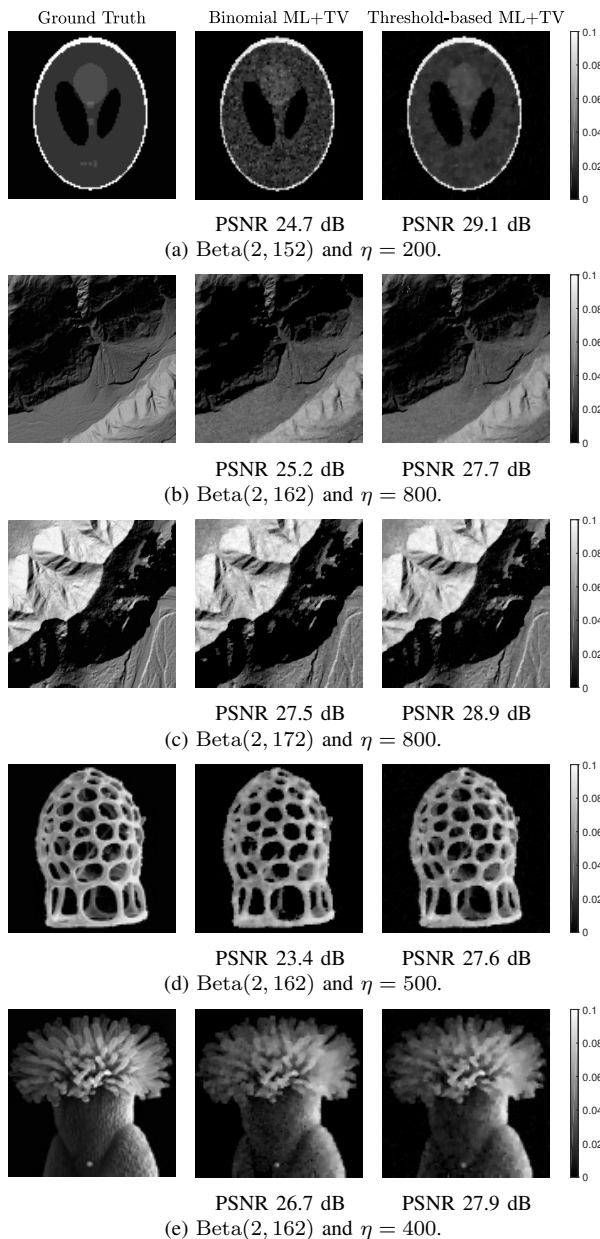


Fig. 15. Images reconstructed through TV-regularized ML estimation showing MSE improvements of up to 4.36 dB for online threshold-based termination in place of conventional binomial sampling. Assumed priors and trial budgets are indicated below each test image. All images are scaled to  $[0.001, 0.101]$ .

## VII. CONCLUSION

We established a novel framework for estimating Bernoulli parameters where we represent each Bernoulli process with a simple trellis graph. By exploiting the mathematical convenience that comes from assuming Beta priors, we propose and study three stopping strategies with varying complexities but yielding very nearly equal performances. All strategies give significant performance improvements over the conventional binomial and negative binomial stopping rules in simulated active imaging applications.

The simple online threshold-based termination was shown to asymptotically allocate trials in the same manner as an oracle-aided solution that assumes the Bernoulli parameters

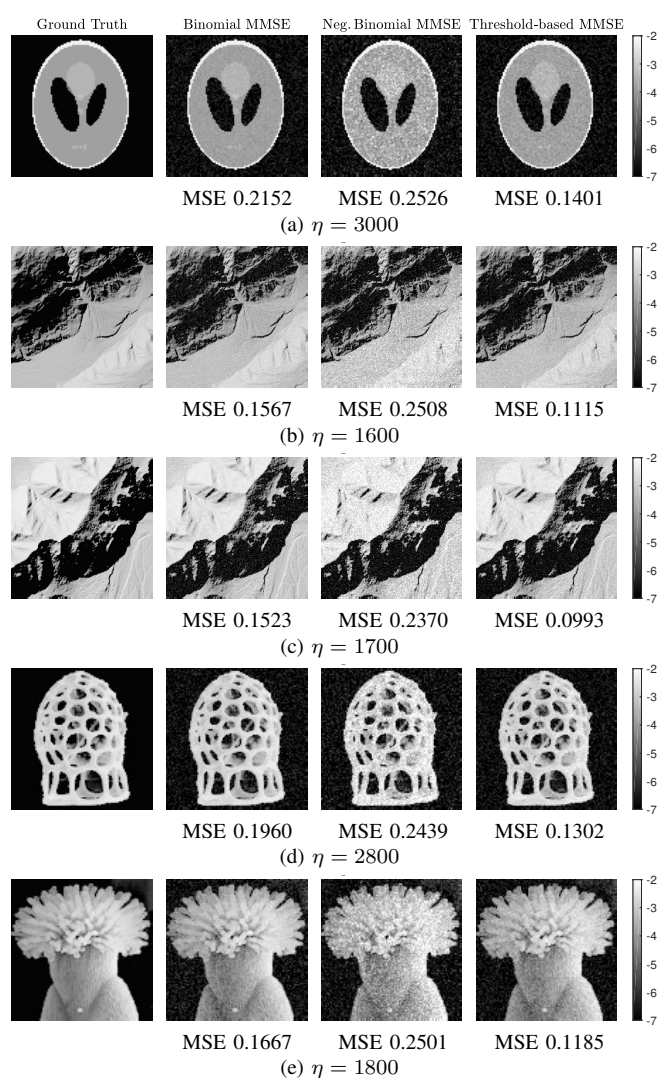


Fig. 16. Estimation of  $\log p$ . Images reconstructed through pixelwise MMSE estimation. All images are scaled to  $[0.001, 0.101]$  and uniform prior is assumed. MSE improvements for online threshold-based termination in place of conventional binomial and negative binomial sampling respectively, are: (a) 1.86 dB and 2.56 dB. (b) 1.48 dB and 3.52 dB. (c) 1.86 dB and 3.78 dB. (d) 1.78 dB and 2.73 dB. (e) 1.48 dB and 3.24 dB. The negative binomial results have been obtained with  $\ell = 5$ .

are *a priori* known. Whilst we only study herein oracle-aided binomial sampling, similar analyses are possible for functions of Bernoulli parameters or oracle-aided negative binomial sampling; these are omitted here because they do not yield clean expressions like the binomial case.

Finally, the proposed online threshold-based termination is extended to the estimation of  $\log p$ . Other functions of  $p$ , such as  $f(p) = 1/p$ , can prove useful for scenarios wherein distinguishing between small parameters is paramount.

In the formulation of optimizing a trellis to minimize MSE, a beta prior is convenient but not at all fundamental. The reduction of the design problem from a general tree to a trellis holds for any prior, and one may in principle compute Bayes risk reduction per trial for any trellis node and any prior. Developing an analogous theory for minimax estimation is also of interest but is less clear because of a lack of additivity

of the cost function. This is especially intriguing because the optimization for MSE incidentally reduces the maximum over  $p$  of the risk (see Fig. 9(b)).

#### ACKNOWLEDGMENT

The authors thank Charles Saunders for assistance with implementation of regularized estimators and Joshua Rapp for discussions on influence in spatial neighborhoods.

#### REFERENCES

- [1] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Photon-efficient computational 3d and reflectivity imaging with single-photon detectors," *IEEE Trans. Comput. Imaging*, vol. 1, pp. 112–125, June 2015.
- [2] M. Wahl, "Time-correlated single photon counting (TCSPC)," tech. rep., PicoQuant, Berlin, Germany, 2014.
- [3] T. L. Fine, *Probability and Probabilistic Reasoning for Electrical Engineering*. Upper Saddle River, NJ: Pearson Prentice Hall, 2006.
- [4] A. Segall, "Bit allocation and encoding for vector sources," *IEEE Trans. Inform. Theory*, vol. IT-22, pp. 162–169, Mar. 1976.
- [5] V. K. Goyal, "Theoretical foundations of transform coding," *IEEE Signal Process. Mag.*, vol. 18, pp. 9–21, Sept. 2001.
- [6] F. J. Anscombe, "Sequential estimation," *J. Roy. Statist. Soc. Ser. B*, vol. 15, no. 1, pp. 1–29, 1953.
- [7] J. B. S. Haldane, "On a method of estimating frequencies," *Biometrika*, vol. 33, pp. 222–225, Nov. 1945.
- [8] M. C. K. Tweedie, "Inverse statistical variates," *Nature*, vol. 155, p. 453, Apr. 14, 1945.
- [9] P. Cabilio and H. Robbins, "Sequential estimation of  $p$  with squared relative error loss," *Proc. Nat. Acad. Sci. USA*, vol. 72, pp. 191–193, Jan. 1975.
- [10] P. Cabilio, "Sequential estimation in Bernoulli trials," *Ann. Statist.*, vol. 5, pp. 342–356, Mar. 1977.
- [11] S. L. Hubert and R. Pyke, "Sequential estimation of functions of  $p$  for Bernoulli trials," in *Game Theory, Optimal Stopping, Probability and Statistics*, vol. 35 of *Lecture Notes-Monograph Series*, pp. 263–294, Institute of Mathematical Statistics, 2000.
- [12] P. M. Djurić and Y. Huang, "Estimation of a Bernoulli parameter  $p$  from imperfect trials," *IEEE Signal Process. Lett.*, vol. 7, pp. 160–163, June 2000.
- [13] D. Ciunzo, A. De Maio, and P. Salvo Rossi, "A systematic framework for composite hypothesis testing of independent Bernoulli trials," *IEEE Signal Process. Lett.*, vol. 22, pp. 1249–1253, Sept. 2015.
- [14] A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [15] J. Rapp and V. K. Goyal, "A few photons among many: Unmixing signal and noise for photon-efficient active imaging," *IEEE Trans. Comput. Imaging*, vol. 3, pp. 445–459, Sept. 2017.
- [16] N. J. Krichel, A. McCarthy, and G. S. Buller, "Resolving range ambiguity in a photon counting depth imager operating at kilometer distances," *Opt. Express*, vol. 18, no. 9, pp. 9192–9206, 2010.
- [17] P. A. Morris, R. S. Aspden, J. E. C. Bell, R. W. Boyd, and M. J. Padgett, "Imaging with a small number of photons," *Nat. Commun.*, vol. 6, Jan. 5, 2015. doi: 10.1038/ncomms6913.
- [18] D. Shin, J. H. Shapiro, and V. K. Goyal, "Single-photon depth imaging using a union-of-subspaces model," *IEEE Signal Process. Lett.*, vol. 22, pp. 2254–2258, Dec. 2015.
- [19] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Lidar waveform-based analysis of depth images constructed using sparse single-photon data," *IEEE Trans. Image Process.*, vol. 25, pp. 1935–1946, May 2016.
- [20] D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. C. Wong, and J. H. Shapiro, "Photon-efficient imaging with a single-photon camera," *Nat. Commun.*, vol. 7, June 24, 2016. doi: 10.1038/ncomms12046.
- [21] D. Shin, J. H. Shapiro, and V. K. Goyal, "Performance analysis of low-flux least-squares single-pixel imaging," *IEEE Signal Process. Lett.*, vol. 23, pp. 1756–1760, Dec. 2016.
- [22] D. Shin, F. Xu, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "Computational multi-depth single-photon imaging," *Opt. Express*, vol. 24, pp. 1873–1888, Feb. 2016.
- [23] L. Mertens, M. Sonnleitner, J. Leach, M. Agnew, and M. J. Padgett, "Image reconstruction from photon sparse data," *Sci. Rep.*, vol. 7, Feb. 7, 2017. doi: 10.1038/srep42164.
- [24] Y. Altmann, R. Aspden, M. Padgett, and S. McLaughlin, "A Bayesian approach to denoising of single-photon binary images," *IEEE Trans. Comput. Imaging*, vol. 3, pp. 460–471, Sept. 2017.
- [25] A. Halimi, A. Maccarone, A. McCarthy, S. McLaughlin, and G. S. Buller, "Object depth profile and reflectivity restoration from sparse single-photon data acquired in underwater environments," *IEEE Trans. Comput. Imaging*, vol. 3, pp. 472–484, Sept. 2017.
- [26] X. Liu, J. Shi, X. Wu, and G. Zeng, "Fast first-photon ghost imaging," *Sci. Rep.*, vol. 8, Mar. 22, 2018. doi: 10.1038/s41598-018-23363-w.
- [27] Y. Altmann, S. McLaughlin, M. J. Padgett, V. K. Goyal, A. O. Hero, and D. Faccio, "Quantum-inspired computational imaging," *Science*, vol. 361, p. 660, Aug. 2018.
- [28] Z. Zhu and S. Pang, "Few-photon computed x-ray imaging," *Appl. Phys. Lett.*, vol. 113, p. 231109, Dec. 2018.
- [29] D. L. Lipke, "Active imaging system using variable gate width time programmed dwell." U.S. Patent 4,151,415, Apr. 1979.
- [30] W. He, Z. Feng, J. Lin, S. Shen, Q. Chen, G. Gu, B. Zhou, and P. Zhang, "Adaptive depth imaging with single-photon detectors," *IEEE Photon. J.*, vol. 9, Apr. 2017.
- [31] T. Dahmen, M. Engstler, C. Pauly, P. Trampert, N. de Jonge, F. Mücklich, and P. Slusallek, "Feature adaptive sampling for scanning electron microscopy," *Sci. Rep.*, vol. 6, May 6, 2016. doi: 10.1038/srep25350.
- [32] S. C. Medin, J. Murray-Bruce, and V. K. Goyal, "Optimal stopping times for estimating Bernoulli parameters with applications to active imaging," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, (Calgary, AB, Canada), pp. 4429–4433, May 2018.
- [33] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1. Belmont, MA, USA: Athena Scientific, 1996.
- [34] R. C. Gonzalez and R. E. Woods, *Digital Image Processing (3rd Edition)*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 2006.
- [35] E. S. Milner and M. T. H. Do, "A population representation of absolute light intensity in the mammalian retina," *Cell*, vol. 171, pp. 865–876.e16, 2018/08/19 2017.
- [36] D. Karlis, "EM algorithm for mixed Poisson and other discrete distributions," *ASTIN Bulletin*, vol. 35, no. 1, p. 324, 2005.
- [37] H. Dette and D. Tomecki, "Hankel determinants of random moment sequences," *J. Theoretical Probability*, vol. 30, pp. 1539–1564, Dec 2017.
- [38] C. Louchet and L. Moisan, "Total variation denoising using posterior expectation," in *Proc. 16th European Signal Process. Conf.*, pp. 1–5, Aug. 2008.
- [39] T. D. Hubbard, M. L. Braun, R. E. Westbrook, and P. E. Gallagher, "High-resolution lidar data for infrastructure corridors, Healy Quadrangle, Alaska." Alaska Division of Geological & Geophysical Surveys, DOI: 10.14509/23163, Dec. 2011.