Optimal Stopping Times for Estimating Bernoulli Parameters with Applications to Active Imaging

Safa C. Medin, John Murray-Bruce, Vivek K Goyal

Boston University Electrical and Computer Engineering Department

Introduction

Estimating the parameter of a **Bernoulli process** p is a fundamental statistical problem with many applications, e.g.:

Photon-efficient active imaging

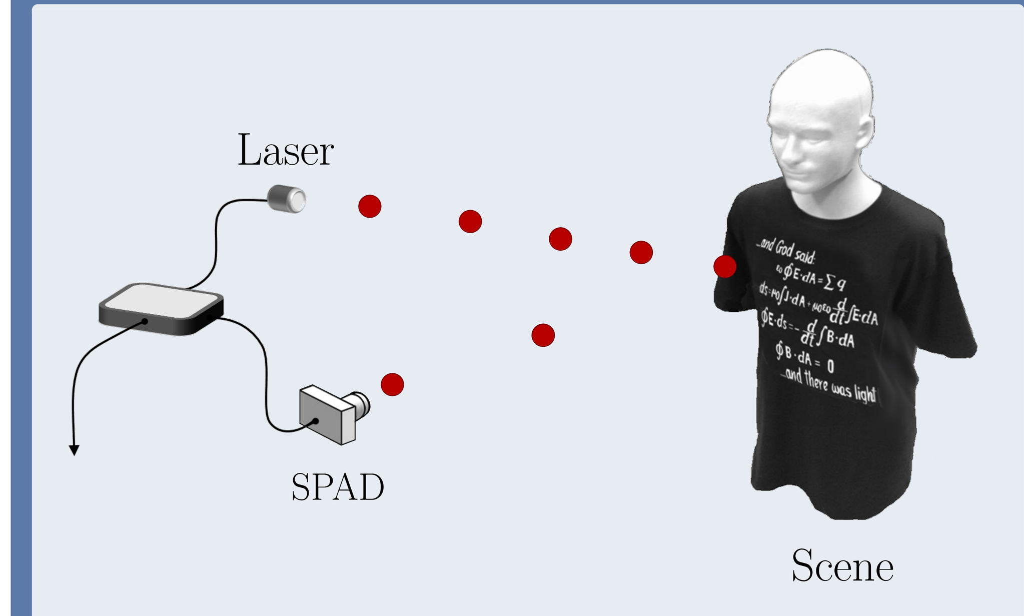

- **Conventional systems:** *nonadaptive*, i.e. number of trials fixed a priori.
- **Alternative system:** data-dependent stopping, known as *sequential estimation*.
- **Motivation:** understanding whether such adaptive systems improve estimation performance.
- **Goal:** Given some trial budget constraint, devise an *optimal stopping strategy* under a mean-squared error loss function.

Contributions

1. Propose a **stopping rule** through a greedy algorithm that seeks for the least achievable error.
2. Generalize stopping rule to a rectangular array of Bernoulli processes, representing pixels in a natural scene.
3. Demonstrate a 4.45 dB improvement in simulated active imaging scenarios.

A Single Bernoulli Process

Probability of continuing observations after trial t:

$$\pi_t : \{0,1\}^t \to [0,1], \qquad t = 0, 1, \ldots$$

Number of observed trials T satisfies $\mathbb{E}[T] \leq n$.

- Any stopping rule can be represented by a sequence of **continuation probabilities.**

Adaptive Stopping Rule

Under Beta(α, β) prior, observing k successes in m trials yields a Beta$(\alpha + k, \beta + m - k)$ distribution.

Bayes risk reduction **from one additional trial**:

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

Proposed data-adaptive stopping rule

Stop when Bayes risk reduction $\Delta R(k, m; \alpha, \beta)$, for an additional trial, is below a specified threshold.

- Visualization of a stopping rule $\to$ **Binary tree** with continuation probability labels.
- All observation sequences with k successes in m trials yield the same continuation probability $q_{k,m} \to$ **Trellis**.

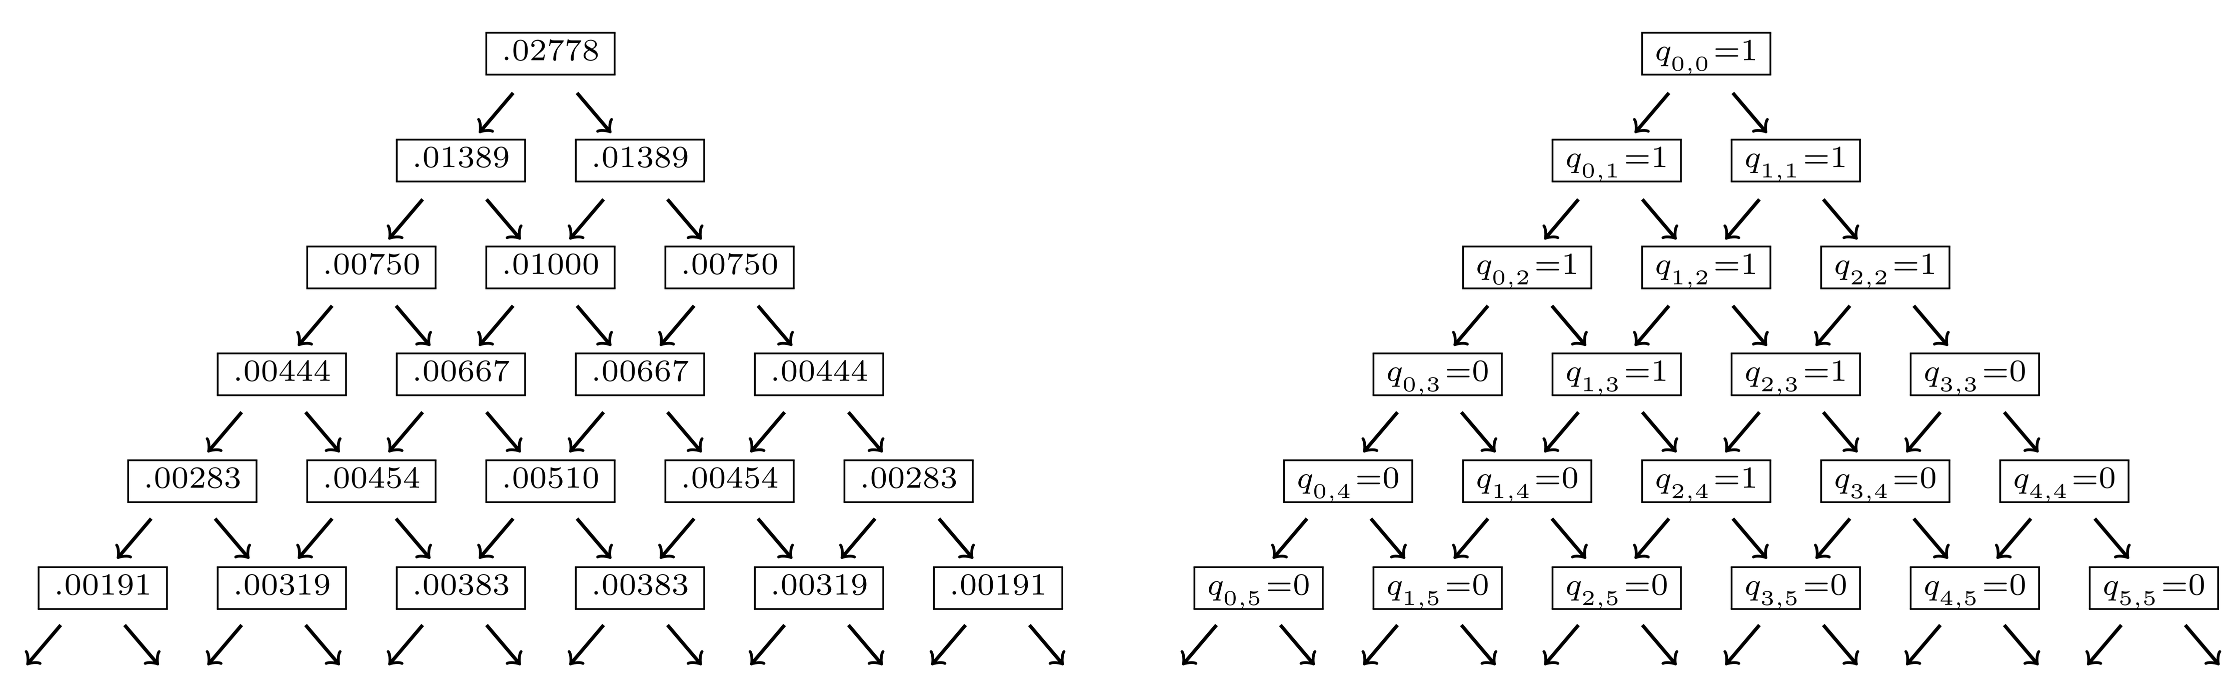

$\Delta R(k, m; \alpha, \beta)$ under Beta$(1, 1)$ (uniform) prior (left) and continuation probabilities for a threshold of 0.005 (right).

- The lower the threshold, the higher the mean number of trials becomes.
- Only certain values of $\mathbb{E}[T]$ are achievable with binary continuation probabilities.
- Small improvement gained for a single Bernoulli process. Byproduct is a significant improvement in imaging applications, due to more efficient allocation of trials across spatial locations.

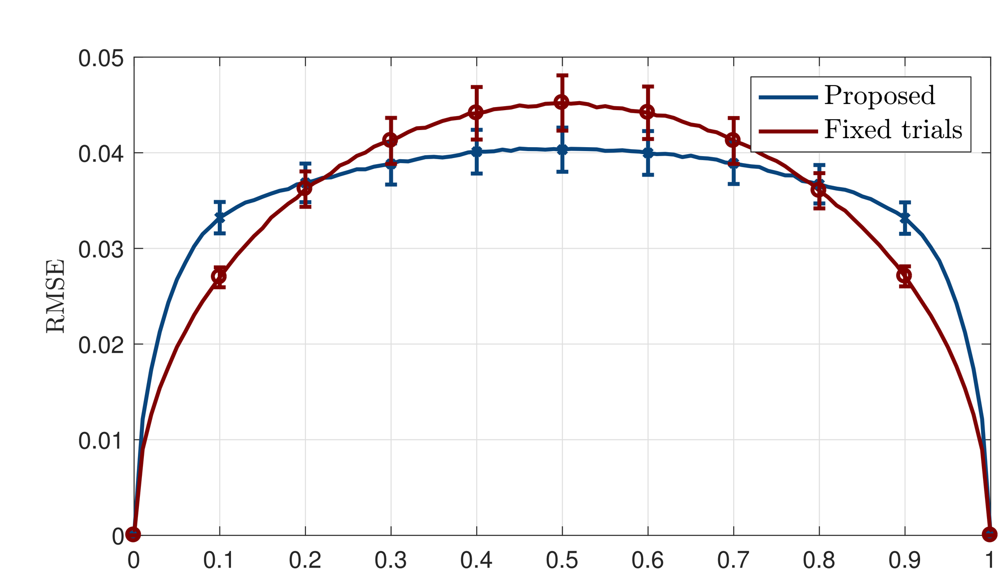

Figure 1: RMSE vs. true Bernoulli parameter, assuming Beta(1,1) prior (uniform) and budget $n = 123$. Mean-squared error reduces from 0.00134 to 0.00129, when averaged over 100 000 experiments.

Arrays of Bernoulli Processes

Scene raster-scanned using pulsed illumination guided by *proposed stopping rule.*

- **Data:** arrays of number of pulses $[m_{i,j}]_{i,j}$ and detections $[k_{i,j}]_{i,j}$.
- **Reconstruction:** TV-regularized ML estimation to exploit spatial correlations.

Results

- True image reflectivity in $[0.001, 0.101]$.
- Beta$(2, 152)$ prior assumed. Trial budget $n = 200$.

Binomial (Fixed) | Adaptive (proposed)

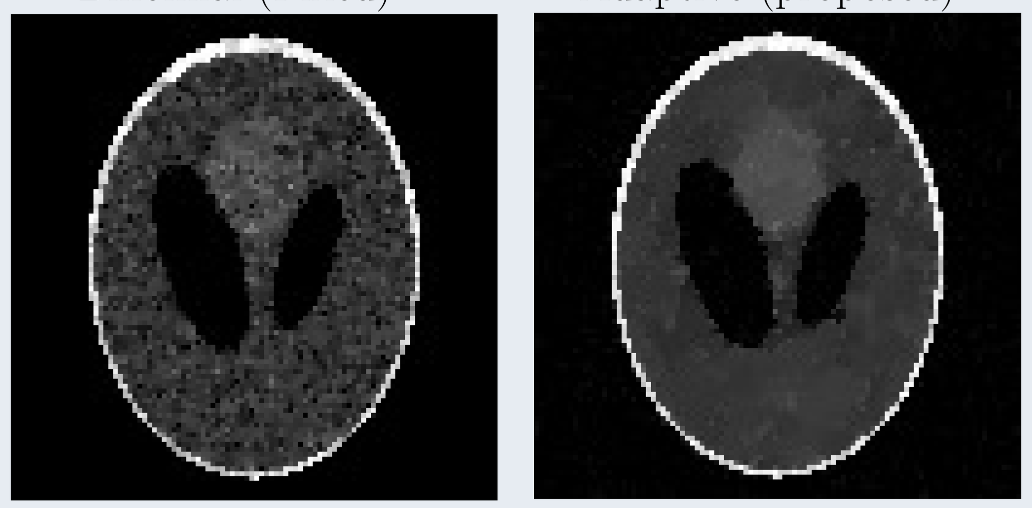

MSE$= 3.4056 \times 10^{-5}$ | MSE$= 1.1779 \times 10^{-5}$

- Average over 100 experiments.

Budget	Method	
	Binomial + TV	Adaptive (proposed) + TV
$n = 58$	9.14e-05	**3.43e-05**
$n = 196$	3.37e-05	**1.26e-05**

Conclusion

- Proposed adaptive stopping rule that yields significant improvements over non-adaptive rule.
- Binomial and Negative Binomial stopping strategies are *rarely* optimal.

References

[1] F. J. Anscombe, 'Sequential estimation,' J. Roy. Statist. Soc. Ser. B, vol. 15, no. 1, pp. 1–29, 1953.

[2] J. B. S. Haldane, "On a method of estimating frequencies,' Biometrika, vol. 33, pp. 222–225, Nov. 1945.

[3] 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.

This material is based upon work supported in part by the US National Science Foundation under Grant No. 1422034.

BOSTON UNIVERSITY