On Computing Optimal Thresholds in Decentralized Sequential Hypothesis Testing

Can Cui and Aditya Mahajan

Electrical and Computer Engineering Department, McGill University

54th Conference on Decision and Control

Outline

Introduction

- Sequential hypothesis testing: sensor network, intrusion detection, primary channel detection, quality control and clinical trials, etc.
- *Decentralized* sequential hypothesis testing: decisions are made in decentralized manner by multi decision makers.
- **Motivation:**
  There are various results that establish optimality of threshold-based strategies in different setups, but few results on how to compute optimal thresholds.

Problem Formulation: Model

Consider a decentralized sequential hypothesis problem investigated in Teneketzis and Ho. (1987).

- **Decision maker:** Two decision makers DM^i, $i \in \{1, 2\}$;
- **Hypothesis:** $H \in \{h_0, h_1\}$ with *a prior* probability p and $1 - p$;
- **Observation:** $Y_t^i \in \mathcal{Y}^i$;
  - $\{Y_t^i\}_{t=1}^{\infty}$ are i.i.d. with PMF f_k^i, $k \in \{1, 2\}$;
  - $\{Y_t^1\}_{t=1}^{\infty}$ and $\{Y_t^2\}_{t=1}^{\infty}$ are conditionally independent given H.
- **Strategy:** $U_t^i \in \{h_0, h_1, \mathsf{C}\}$ according to $U_t^i = g_t^i(Y_{1:t}^i)$.

Problem Formulation: Model

- **Stopping time:** $N^i = \min\{t \in \mathbb{Z} > 0 : U_t^i \in \{h_0, h_1\}\}$;
- **Observation cost:** c^i for each observation at DM^i;
- **Stopping cost:** $\ell(U^1, U^2, H)$ which satisfies:
  - $\ell(U^1, U^2, H)$ cannot be decomposed as $\ell(U^1, H) + \ell(U^2, H)$;
  - For any $m, n \in \{h_0, h_1\}$, $m \neq n$,

$$\ell(m, m, n) \geqslant \ell(n, m, n) \geqslant c^i \geqslant \ell(n, n, n);$$
$$\ell(m, m, n) \geqslant \ell(m, n, n) \geqslant c^i \geqslant \ell(n, n, n).$$

- **Goal:** Given p, choose (g^1, g^2) to minimize $J(g^1, g^2; p)$, where

$$J(g^1, g^2; p) = \mathbb{E}[c^1 N^1 + c^2 N^2 + \ell(U^1, U^2, H)].$$

Problem Formulation: Problem

- **Problem 1**: Given the prior probability p, the observation PMFs f_0^i, f_1^i, the observation cost c^i, and the loss function ℓ, find a strategy (g^1, g^2) that minimizes the cost given by$J(g^1, g^2; p)$.

- **Problem 2**: Given the prior probability p, the observation PMFs f_0^i, f_1^i, the observation cost c^i, and the loss function ℓ, find a strategy (g^1, g^2) that is person-by-person optimal (PBPO).

  A person-by-person optimal (PBPO) strategy (g^1, g^2) satisfies:

$$J(g^1, g^2) \leq J(g^1, \tilde{g}^2), \quad \forall \tilde{g}^2 \in \mathcal{G}^2,$$
$$J(g^1, g^2) \leq J(\tilde{g}^1, g^2), \quad \forall \tilde{g}^1 \in \mathcal{G}^1.$$

Information State Process

For any $i \in \{1,2\}$, let $-i$ denote the other decision maker. For any realization $y^i_{1:t}$ of $Y^i_{1:t}$, define

$$\pi^i_t := \mathbb{P}(H = h_0 \mid y^i_{1:t}).$$

In addition, define

$$q^i(y^i_{t+1} \mid \pi^i_t) := \pi^i_t \cdot f^i_0(y^i_{t+1}) + (1 - \pi^i_t) \cdot f^i_1(y^i_{t+1}), \tag{1}$$

$$\phi^i(\pi^i_t, y^i_{t+1}) := \pi^i_t \cdot f^i_0(y^i_{t+1}) / q^i(y^i_{t+1} \mid \pi^i_t). \tag{2}$$

The update of the information state is given by $\pi^i_{t+1} = \phi^i(\pi^i_t, y^i_{t+1})$.

$\{\pi^i_t\}_{t=1}^{\infty}$ is an information state process for DM^i.

For ease of notation, for any $i \in \{1,2\}$, $k \in \{0,1\}$, $u^i \in \{h_0, h_1\}$, and $g^i \in \mathcal{G}^i$, define

$$\xi^i_k(u^i, g^i; p) = \mathbb{P}(U^i = u^i \mid H = h_k; g^i, p).$$

Structure of Optimal Decision Rules

Threshold based strategy: A strategy of the above form is called threshold based if there exists thresholds $\alpha_t^i, \beta_t^i \in [0, 1]$, $\alpha_t^i \leq \beta_t^i$, such that for any $\pi^i \in [0, 1]$,

$$g_t^i(\pi^i) = \begin{cases} h_1 & \text{if } \pi^i < \alpha_t^i, \\ \mathsf{C} & \text{if } \alpha_t^i \leq \pi^i \leq \beta_t^i, \\ h_0 & \text{if } \pi^i > \beta_t^i. \end{cases}$$

Time invariant strategy: A strategy $g^i = (g_1^i, g_2^i, \dots)$ is called time invariant if for any $\pi^i \in [0, 1]$, $g_t^i(\pi^i)$ does not depend on t.

Theorem

For any $i \in \{1, 2\}$ and any time-invariant and threshold-based strategy $g^{-i} \in \mathcal{G}^{-i}$, there is no loss of optimality in restricting attention to time-invariant and threshold based strategies at DM^i. Moreover, the best response strategy at DM^i is given by the solution of a dynamic program:

Dynamic Program

For any $\pi^i \in [0, 1]$

$$V^i(\pi^i) = \min\{W_0^i(\pi^i, g^{-i}), W_1^i(\pi^i, g^{-i}), W_{\mathsf{C}}^i(\pi^i, g^{-i})\}, \tag{3}$$

where for $k \in \{0, 1\}$,

$$W_k^1(\pi^1, g^2) = \sum_{u^2 \in \{h_0, h_1\}} \big[\xi_0^2(u^2, g^2; \pi^1) \cdot \pi^1 \cdot \ell(h_k, u^2, h_0) + \xi_1^2(u^2, g^2; \pi^1) \cdot (1 - \pi^1) \cdot \ell(h_k, u^2, h_1)\big], \tag{4}$$

W_k^2 is defined similarly, and

$$W_{\mathsf{C}}^i(\pi^i, g^{-i}) = c^i + \mathscr{B}^i V^i(\pi^i), \tag{5}$$

where $\mathscr{B}^i$ is the Bellman operator given by

$$[\mathscr{B}^i V^i](\pi^i) = \sum_{y^i} q(y^i \mid \pi^i) \cdot V^i(\phi(\pi^i, y^i)),$$

and $q(y^i \mid \pi^i)$ is given by (**??**).

Algorithms for computing optimal thresholds

We propose two methods to compute the optimal thresholds.

- Orthogonal search
  Iteratively solve

$$\langle \alpha^1, \beta^1 \rangle = \mathcal{D}^1(\langle \alpha^2, \beta^2 \rangle) \quad \text{and} \quad \langle \alpha^2, \beta^2 \rangle = \mathcal{D}^2(\langle \alpha^1, \beta^1 \rangle). \tag{6}$$

- Direct search
  Approximately compute $J(\langle \alpha^1, \beta^1 \rangle, \langle \alpha^2, \beta^2 \rangle; p)$ and search for optimal $\langle \alpha^1, \beta^1 \rangle$, $\langle \alpha^2, \beta^2 \rangle$ using derivative-free non-convex optimization method.

Orthogonal search

The following procedures are used to solve the coupled dynamic programs:

1. Start with an arbitrary threshold-based strategy $(\langle \alpha^1_{(1)}, \beta^1_{(1)} \rangle)$.
2. Construct a sequence of strategies as follows:
   1. For even n:
      $$\langle \alpha^1_{(n)}, \beta^1_{(n)} \rangle = \mathcal{D}^1(\langle \alpha^2_{(n-1)}, \beta^2_{(n-1)} \rangle),$$
      and
      $$\langle \alpha^2_{(n)}, \beta^2_{(n)} \rangle = \langle \alpha^2_{(n-1)}, \beta^2_{(n-1)} \rangle.$$
   2. For odd n:
      $$\langle \alpha^1_{(n)}, \beta^1_{(n)} \rangle = \langle \alpha^1_{(n-1)}, \beta^1_{(n-1)} \rangle,$$
      and
      $$\langle \alpha^2_{(n)}, \beta^2_{(n)} \rangle = \mathcal{D}^2(\langle \alpha^1_{(n-1)}, \beta^1_{(n-1)} \rangle).$$

Orthogonal search

Theorem

The orthogonal search procedure described above converges to a time-invariant threshold-based strategy (g^1, g^2) that is person-by-person optimal.

Proof.

Let $(g^1_{(n)}, g^2_{(n)})$ denote the strategy at step n. By construction,

$$J(g^1_{(n)}, g^2_{(n)}) \leq J(g^1_{(n-1)}, g^2_{(n-1)}).$$

Thus, the sequence $\{J(g^1_{(n)}, g^2_{(n)})\}$ is a decreasing sequence lower bounded by 0. Hence, a limit exists and the limiting strategy is PBPO. □

Preliminaries: Discretizing continuous state Markov chains

For any $m \in \mathbb{N}$, for any $i \in \{0, 1\}$, we approximate the $[0, 1]$-valued Markov process $\{\pi_t^i\}_{t=1}^{\infty}$, by a $\mathcal{S}_m$-valued Markov chain $\mathcal{S}_m = \{0, \frac{1}{m}, \frac{2}{m}, \ldots, 1\}$.

Algorithm 1: Compute transition matrices

input: Discretization size m, DM i; **output**: P_0^i, P_1^i, P_*^i

forall $s_p \in \mathcal{S}_m$ **do**

- **forall** $y \in \mathcal{Y}^i$ **do**
  - let $s_+ = \phi^i(s, y^i)$
  - find $s_q, s_{q+1} \in \mathcal{S}_m$ such that $s_+ \in [s_q, s_{q+1})$
  - find $\lambda_q^y, \lambda_{q+1}^y \in [0, 1]$ such that
    - $\lambda_q^y + \lambda_{q+1}^y = 1$
    - $s_+ = \lambda_q^y s_q + \lambda_{q+1}^y s_{q+1}$
- **forall** $q \in \{0, 1, \ldots, m\}$ **do**
  - $[P_0^i]_{pq} = \sum_y \lambda_q^y \cdot f_0^i(y) \cdot s_p$
  - $[P_1^i]_{pq} = \sum_y \lambda_q^y \cdot f_1^i(y) \cdot (1 - s_p)$
  - $[P_*^i]_{pq} = \sum_y \lambda_q^y \cdot q^i(y^i \mid s_p)$

Approximation with discrete-state Markov chain

For a given ξ^i_k (fix i and k), given any threshold based strategy $g^i = \langle \alpha^i, \beta^i \rangle$ such that $\alpha^i, \beta^i \in \mathcal{S}_m$, define sets $\mathcal{A}^i_0, \mathcal{A}^i_1 \subset \mathcal{S}_m$ as: $\mathcal{A}^i_0 = \left\{\beta^i, \beta^i + \frac{1}{m}, \dots, 1\right\}$ and $\mathcal{A}^i_1 = \left\{0, \frac{1}{m}, \dots, \alpha^i\right\}$ as shown below.

Then $\xi^i_k(h_0, g^i; p)$ is approximated by the event that the Markov chain with transition probability P^i_k that starts in p gets absorbed in the set $\mathcal{A}^i_0$ before it is absorbed in the set $\mathcal{A}^i_1$.

Define $\theta^i_k(g^i; p) = \mathbb{E}[N_i \mid H = h_k; g^i, p]$, then $\theta^i_k(g^i; p)$ can be approximated using the expected stopping time of Markov chain. This is approximated by the event that the Markov chain starting in p is absorbed in $(\mathcal{A}^i_b \cup \mathcal{A}^i_1)$.

Approximation with discrete-state Markov chain

Let $\hat{P}_k^i$ be the transition matrix of the corresponding absorbing Markov chain. Re-order states so that $\hat{P}_k^i$ may be written in the canonical form

$$\hat{P}_k^i = \begin{pmatrix} Q_k^i & R_k^i \\ 0 & I \end{pmatrix},$$

Define $B_k^i = (1 - Q_k^i)^{-1} R_k^i$, then,

$$\xi_k^i(h_b, \langle \alpha^i, \beta^i \rangle; p) \approx [B_k^i]_{p^* b}, \quad b \in \{0, 1\}, \tag{7}$$

Define $T_k^i = (I - Q_k^i)^{-1} \mathbf{1}$, where $\mathbf{1}$ is a column vector with all entries as 1, then,

$$\theta_k^i(\langle \alpha^i, \beta^i \rangle; p) \approx [T_k^i]_{p^*}, \tag{8}$$

where p^* denotes the index of p in $\mathcal{S}_m \setminus (\mathcal{A}_0^i \cup \mathcal{A}_1^i)$.

Approximate solution of the dynamic program

We've approximated $\xi^i_k(\cdot, g^{-i}; p)$, and therefore approximately compute $W^i_k(\pi^i, g^{-i})$.

Define an approximate Bellman operator using the first-order hold transition matrix P^i_* as follows:

$$[\hat{\mathscr{B}}^i V^i](s) = c^i + \sum_{s_+ \in \mathcal{S}_m} [P^i_*]_{ss_+} V(s_+).$$

Then $\hat{\mathscr{B}}^i$ corresponds to the discretization of $\mathscr{B}^i$ on $\mathcal{S}_m$ and performing linear interpolation on points outside $\mathcal{S}_m$. Hence, it may be used to approximately compute $W_{\mathsf{C}}(\pi^i, g^{-i})$.

Combing all these, we get an approximate procedure to solve the dynamic program of Theorem 1. This, in turn, gives an approximate procedure for finding a PBPO strategy using orthogonal search.

Direct search: Performance of an arbitrary strategy

Recall the definition and approximation of $\xi_k^i(u^i, g^i; p)$ and $\theta_k^i(g^i; p)$. For a particular *a prior* probability p, the expected cost $J(g^1, g^2; p)$ can be expanded as:

$$
\begin{aligned}
J(g^1, g^2; p) &= p \cdot [c^1 \cdot \theta_0^1(g^1; p) + c^2 \cdot \theta_0^2(g^2; p)] \\
&\quad + (1-p) \cdot [c^1 \cdot \theta_1^1(g^1; p) + c^2 \cdot \theta_1^2(g^2; p)] \\
&\quad + \sum_{u^1, u^2 \in \{h_0, h_1\}}^{2} \Big[p \cdot \xi_0^1(u^1, g^1; p) \cdot \xi_0^2(u^2, g^2; p) \cdot \ell(u^1, u^2, h_0) \\
&\qquad + (1-p) \cdot \xi_1^1(u^1, g^1; p) \cdot \xi_1^2(u^2, g^2; p) \cdot \ell(u^1, u^2, h_1) \Big]. \qquad (9)
\end{aligned}
$$

Direct search: Search over all threshold based strategy.

We expect $J(p, \langle\alpha^1, \beta^1\rangle, \langle\alpha^2, \beta^2\rangle)$ to be non-convex in the parameters $(\alpha^1, \beta^1, \alpha^2, \beta^2)$. Since there is no analytic expression for J, in the numerical results We use a derivative-free algorithms—Nelder-Mead simplex algorithm.

To reduce the dependence of the numerical results on the choice of the *a priori* probability p, we pick multiple values of p in a finite set $\mathcal{P} \subset [0, 1]$ and use

$$\hat{J}(\alpha^1, \beta^1, \alpha^2, \beta^2) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} J(p, \langle\alpha^1, \beta^1\rangle, \langle\alpha^2, \beta^2\rangle)$$

as the objective function for the non-convex optimization algorithm. If $J(p, \langle\alpha^1, \beta^1\rangle, \langle\alpha^2, \beta^2\rangle)$ was computed exactly, then such an averaging will not affect the result of the optimization algorithm because the optimal strategy (g^1, g^2) does not depend on the choice of p.

We compare the performance of orthogonal search and direct search on a benchmark system, $\mathcal{Y}^1 = \mathcal{Y}^2 = \{0, 1\}$ and the loss function is of the form:

$$\ell(u^1, u^2, h) = \begin{cases} 0, & \text{if } u^1 = u^2 = h, \\ 1, & \text{if } u^1 \neq u^2, \\ L, & \text{if } u^1 = u^2 \neq h. \end{cases} \tag{10}$$

For both methods, we use $m = 1000$ and in direct search, we use $\mathcal{P} = \mathcal{S}_m$.

Note that by choosing parameters (c^1, c^2, L) and observation distributions $(f_0^1, f_1^1, f_0^2, f_1^2)$, we completely specifies the model.

We will work with two choices of parameters:

- A particular instance.
  A system with $c^1 = c^2 = 0.05$, and

$$f_0^1 = \begin{bmatrix} 0.25 & 0.75 \end{bmatrix}, \qquad f_0^2 = \begin{bmatrix} 0.80 & 0.20 \end{bmatrix},$$
$$f_1^1 = \begin{bmatrix} 0.60 & 0.40 \end{bmatrix}, \qquad f_0^2 = \begin{bmatrix} 0.30 & 0.70 \end{bmatrix}.$$

  In coupled loss cases $L = 2.5$, in decomposables cases $L = 2$.
- Randomized parameters.
  Randomly generate 500 instances of the parameters (c^1, c^2, L) and $(f_0^1, f_1^1, f_0^2, f_1^2)$. Specifically, we use $f_k^i = [\delta_k^i, 1 - \delta_k^i]$ with $\delta_k^i \sim \text{unif}[0, 1]$.
  In decomplsable cases, $L = 2$

In the following slides, we will show the numerical results for three scenarios based on the parameters described above.

Coupled Loss Case

Let OS and DS denote the solution obtained by orthogonal search and direct search.

	$g^1 = \langle \alpha^1, \beta^1 \rangle$	$g^2 = \langle \alpha^2, \beta^2 \rangle$	$\hat{J}(g^1, g^2)$	iters.	runtime
OS	$\langle 0.326, 0.73 \rangle$	$\langle 0.07, 0.931 \rangle$	0.455	5	1.45s
DS	$\langle 0.287, 0.726 \rangle$	$\langle 0.14, 0.863 \rangle$	0.436	45	6.05s

Let J_{OS}, J_{DS} denote the performance of the solution obtained by orthogonal search and direct search. Define $\Delta J_{OS} = (J_{OS} - J_{DS})/J_{OS}$ and $\Delta J_{DS} = (J_{DS} - J_{OS})/J_{DS}$.

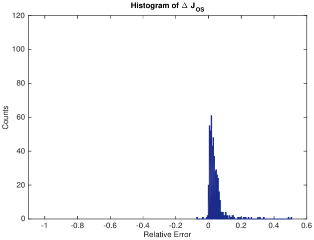

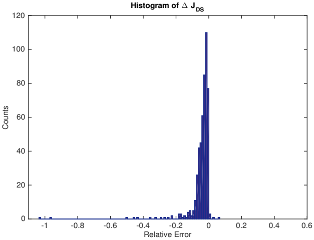

Decomposable Case

The problem decomposes into two centralized problem when $\ell(U^1, U^2, H)$ equals to $\ell(U^1, H) + \ell(U^2, H)$. We use value iteration to solve centralized problem and refer to this solution as *centralized solution*, denoted as CS.

	$g^1 = \langle \alpha^1, \beta^1 \rangle$	$g^2 = \langle \alpha^2, \beta^2 \rangle$	$\hat{J}(g^1, g^2)$
OS	$\langle 0.318, 0.686 \rangle$	$\langle 0.089, 0.913 \rangle$	0.428
DS	$\langle 0.3053, 0.7055 \rangle$	$\langle 0.1845, 0.8218 \rangle$	0.406
CS	$\langle 0.305, 0.705 \rangle$	$\langle 0.184, 0.822 \rangle$	0.406

Let J^* denote the centralized solution. Define the relative errors $E_{OS} = (J_{OS} - J^*)/J^*$ and $E_{DS} = (J_{DS} - J^*)/J^*$.

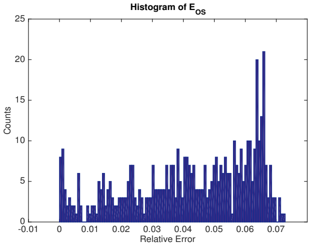

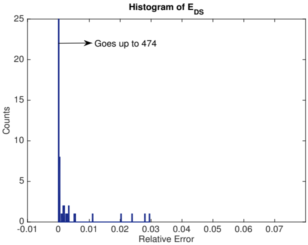

Asymptotic Case

When $c^i \ll L$, the asymptotic expression of ξ_k^i and θ_k^i are given as:

$$\xi_0^i(h_1, g^i, p) = \frac{\alpha^i(1-p)}{(1-\alpha^i)p} = B, \quad \xi_1^i(h_0, g^i, p) = \frac{(1-\beta^i)p}{\beta^i(1-p)} = A.$$

$$\theta_0^i(p, g^i) = \frac{\log(A)}{\sum_{Y^i}[\log\frac{f_1^i(Y^i)}{f_0^i(Y^i)}]\cdot f_0^i(Y^i)}, \quad \theta_1^i(p, g^i) = \frac{\log(1/B)}{\sum_{Y^i}[\log\frac{f_1^i(Y^i)}{f_0^i(Y^i)}]\cdot f_0^i(Y^i)}.$$

Then use direct search to find the optimal threshold g^i. The histograms of E_{OS} and E_{DS} are shown below.

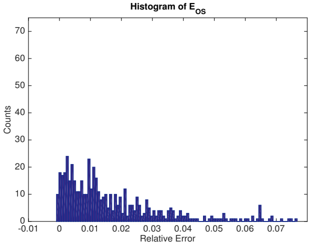

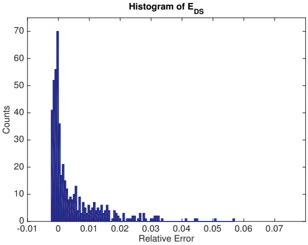

Summary

- Two methods to approximately compute the optimal threshold-based strategies in decentralized sequential hypothesis testing.
- Discretization of continuous-valued information state process by a finite-valued Markov chain.
- In our example, direct search performs better than orthogonal search; sometimes, significantly better.

A future direction is to generalize the approximation methods developed in this paper to more general decentralized sequential hypothesis models.