Active Sequential Hypothesis Testing with Application to a Visual Search Problem

Nidhin Koshy Vaidhiyan
nidhinkv@ece.iisc.ernet.in
Department of ECE
Indian Institute of Science
Bangalore 560012, India

S. P. Arun
sparun@cns.iisc.ernet.in
Centre for Neuroscience
Indian Institute of Science
Bangalore 560012, India

Rajesh Sundaresan
rajeshs@ece.iisc.ernet.in
Department of ECE
Indian Institute of Science
Bangalore 560012, India

***Abstract*—We consider a visual search problem studied by Sripati and Olson where the objective is to identify an oddball image embedded among multiple distractor images as quickly as possible. We model this visual search task as an active sequential hypothesis testing problem (ASHT problem). Chernoff in 1959 proposed a policy in which the expected delay to decision is asymptotically optimal. The asymptotics is under vanishing error probabilities. We first prove a stronger property on the moments of the delay until a decision, under the same asymptotics. Applying the result to the visual search problem, we then propose a "neuronal metric" on the measured neuronal responses that captures the discriminability between images. From empirical study we obtain a remarkable correlation ($r = 0.90$) between the proposed neuronal metric and speed of discrimination between the images. Although this correlation is lower than with the L_1 metric used by Sripati and Olson, this metric has the advantage of being firmly grounded in formal decision theory.**

I. Introduction

Sripati and Olson [1] studied the correlation between the speed of discrimination of images and a measure of dissimilarity of neuronal representations of the images in the inferotemporal (IT) cortex of the brain. They conducted two sets of experiments, one on human subjects and the other on macaque monkeys. The experiments were the following.
(1) In the experiment on humans, subjects were shown a picture as in Figure 1(a) having six images placed at the vertices of a regular hexagon, with one image being different from the others. In particular, let I_1 and I_2 be two images. One of these two was picked randomly with equal probability and was placed at one of the six locations randomly, again with equal probability. The remaining five locations contained the copies of the other image. See Figure 1(a) and 1(b). The subjects were required to identify the correct half (left or right) of the plane where the odd image was located. The subjects were advised to indicate their decision "as quickly as possible without guessing" [1]. The time taken to make a decision after the onset of the image was recorded.
(2) In the other set of experiments on macaque monkeys, the images I_1 and I_2 were displayed on the screen, and the neuronal firings elicited by I_1 separately and I_2 separately on a set of IT neurons were recorded. The neuronal representation of an image was taken to be a vector of average firing rates across the neurons. The two experiments were done on several pairs (I_1,I_2) of images. The authors observed a remarkably high correlation ($r = 0.95$) between the speed of discrimination in humans and the L_1 distance between the firing rate vectors of the two images. They concluded that the more similar the neuronal representation (in monkey IT cortex), the tougher it is for humans to distinguish the images.

In this paper, we model the visual search problem as an active sequential multiple hypothesis testing problem (ASHT problem), first studied by Chernoff [2]. In sequential hypothesis testing [3]–[9], the only decision a controller makes at each stage is to either stop and make a decision or continue to draw another sample and trade off the cost of delay for a decision of better quality. In active sequential hypothesis testing, the controller is additionally capable of controlling the quality of the sample drawn. In our experiment the human subject can actively choose the location to sample based on past observations and actions.

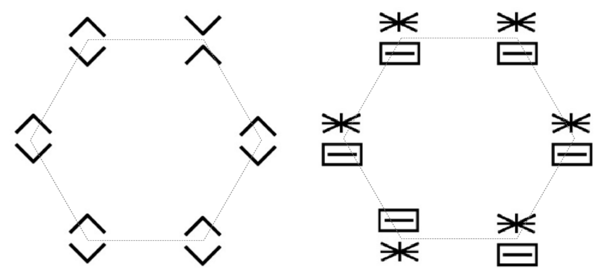

(a) (b)

Fig. 1. Two examples of search task configurations.

a) Prior Work: Chernoff [2] formulated the ASHT problem, proved an asymptotic lower bound under vanishing error probabilities and further devised the so-called *Procedure A* that he went on to show was asymptotically optimal under vanishing error probabilities. Recently Naghshvar and Javidi [10]–[12], also motivated by a visual search problem, cast Chernoff's problem into the framework of dynamic programming and attempted to derive structural properties on the value functions and the decision regions. A related but different line of research was carried out by Rao [13] who studied visual search as a partially observed Markov decision problem over fixed time horizons. In contrast, the works of Chernoff [2], Naghshvar and Javidi [10]–[12], and this work model the

978-1-4673-2579-0/12/$31.00 ©2012 IEEE

search problem as one of optimal stopping with controls.

b) Our contribution: We first provide a mild extension of Chernoff's asymptotic optimality of expected delay for his *Procedure A* by showing asymptotic optimality for all moments. We then recognize that Chernoff's solution suggests a natural "neuronal metric" for the visual search problem that is more appropriate than the L_1 distance between the vectors of firing rates used in Sripati and Olson [1]. This metric is closely related to the relative entropy between the Poisson point processes associated with firing rate vectors for the two images. We do see a high correlation ($r = 0.9$) between speed of discrimination and the new neuronal metric. While this is lower than the correlation ($r = 0.95$) with the L_1 metric as observed by Sripati and Olson [1], we remark that our abstraction is overly simplistic, and emphasize that our work provides a more logical basis for the proposed neuronal metric than the L_1 distance.

c) Organization: In Section II we study the ASHT problem. We first set up the notation, state the assumptions, and identify the performance criterion in sections II-A through II-C. In section II-D we give an asymptotic lower bound for the moments of the stopping time for all policies that satisfy a constraint on the maximum possible conditional error cost. In section II-E we show that a class of policies proposed by Chernoff [2] and Naghshvar and Javidi [12] asymptotically attain the lower bounds for general moments. In Section III, we return to the visual search problem, and obtain the appropriate metric for comparison with delays and identify the degree of correlation.

II. The ASHT Problem

In this section we study the active sequential multiple hypothesis testing problem (ASHT problem).

A. Basic Notation

Let H_i, $i = 1, 2, \ldots, M$ denote the M hypotheses of which only one holds true. Let $\mathcal{A}$ be the set of all possible actions which we take as finite $|\mathcal{A}| = K < \infty$. Let $\mathcal{X}$ be the observation space. Let $(X_n)_{n\geq 1}$ and $(A_n)_{n\geq 1}$ denote the observation process and the control process respectively. We assume that $(A_n)_{n\geq 1}$ is a random or deterministic function of the past observations and actions. Conditioned on action A_n and the true hypothesis H, we assume that X_n is conditionally independent of previous actions $\mathbf{A^{n-1}} = (A_1, A_2, \ldots A_{n-1})$ and observations $\mathbf{X^{n-1}} = (X_1, X_2, \ldots, X_{n-1})$. Let q_i^a be the probability density function with respect to some reference measure μ for observation X under action a when H_i is the true hypothesis. Let $q_i(\mathbf{X^n}, \mathbf{A^n})$ be the probability density, with respect to a common reference measure $\mu^{\otimes n} \times \mathsf{unif}(\mathcal{A})^{\otimes n}$, of observations and actions till time n, where $\mathsf{unif}(\mathcal{A})$ is the uniform distribution on $\mathcal{A}$. Let $Z_{ij}(n)$ denote the log-likelihood ratio (LLR) process of hypothesis H_i with respect to hypothesis H_j, i.e.,

$$Z_{ij}(n) = \log \frac{q_i\left(\mathbf{X^n}, \mathbf{A^n}\right)}{q_j\left(\mathbf{X^n}, \mathbf{A^n}\right)} = \sum_{l=1}^{n} \log \frac{q_i^{A_l}\left(X_l\right)}{q_j^{A_l}\left(X_l\right)}. \quad (1)$$

Let E_i denote the conditional expectation and let P_i denote the conditional probability measure under the hypothesis H_i. Let $D(q_i^a \| q_j^a)$ denote the relative entropy between the probability measures associated with the observations under hypothesis H_i and hypothesis H_j under action a.

Let $p_i(0)$ denote the prior probability that hypothesis H_i is true. The posterior probability that hypothesis H_i is true, given the observations and actions till time n is denoted $p_i(n)$. The beliefs $\mathbf{p}(n) = (p_1(n), \ldots, p_M(n))$ admit the sequential update

$$p_i(n+1) = \frac{p_i(n) q_i^{A_{n+1}}(X_{n+1})}{\sum_j p_j(n) q_j^{A_{n+1}}(X_{n+1})}, \quad (2)$$

a fact that follows by an easy application of Bayes rule.

A policy π is a sequence of action plans that at time n looks at the history $\mathbf{X}^{n-1}, \mathbf{A}^{n-1}$ and prescribes a composite action that could be either $(stop, d)$ or $(continue, \lambda)$. If the composite action is $(stop, d)$, then d is the decision on the hypothesis at the stopping time, and so $d \in \{1, 2, \ldots, M\}$. If the composite action plan is $(continue, \lambda)$, then $\lambda \in \mathcal{P}(\mathcal{A})$ is the distribution with which the next control is picked. Let τ be the stopping time. Recall that $(A_n)_{n\geq 1}$ is the control process until the stopping time.

Given an error tolerance vector $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_M)$ with $0 < \alpha_i < 1$, let $\Delta(\alpha)$ be the set of policies

$$\Delta(\alpha) = \left\{\pi : P_i(d \neq i) \leq \alpha_i, \ \forall\, i\right\}. \quad (3)$$

These are policies that meet a specified conditional error probability tolerance criterion. Let $||\alpha|| = \max_i \alpha_i$.

We define λ_i to be the mixed action under hypothesis H_i that guards against the nearest alternative, i.e., $\lambda_i \in \mathcal{P}(\mathcal{A})$ such that

$$\lambda_i := \arg\max_{\lambda} \min_{j\neq i} \sum_{a} \lambda_i(a) D(q_i^a \| q_j^a). \quad (4)$$

Further define

$$D_i = \max_{\lambda} \min_{j\neq i} \sum_{a\in\mathcal{A}} \lambda(a) D\left(q_i^a \| q_j^a\right) \quad (5)$$

Let $\mathcal{A}_{ij} = \{a \in \mathcal{A} : D(q_i^a \| q_j^a) > 0\}$ be the set of all actions that can differentiate hypothesis H_i from hypothesis H_j. It is easy to see that $\mathcal{A}_{ij} = \mathcal{A}_{ji}$. Finally, let $\beta_{ij}^k = \sum_{a\in\mathcal{A}_{ij}} \lambda_k(a)$, i.e., β_{ij}^k is the probability of choosing some action that can distinguish hypothesis H_i from hypothesis H_j when the actions are picked according to λ_k.

B. Assumptions

Throughout, we assume the following:

(I) $E_i\left[\left(\log \frac{q_i^a(X)}{q_j^a(X)}\right)^2\right] < \infty \ \forall\, i, j, a.$

(II) $\beta = \min_{i,j,k} \beta_{ij}^k > 0.$

Assumption (I) implies that $D(q_i^a || q_j^a) < \infty$ and ensures that no single observation can result in a reliable decision. Assumption (II) is an important technical assumption in our work. A crucial exponential boundedness property of the stopping time for Chernoffs *Procedure A* and Naghshvar and Javidi's Policy-NJ (described in section II-E), which is required in many proofs, is based on this assumption. Further it, ensures that under all beliefs, there is a positive probability of choosing an action that can distinguish hypothesis H_i from hypothesis H_j $(i \neq j)$. In particular, for any distinct i and j, there is at least one control that can help distinguish the hypotheses H_i from H_j.

C. Performance Criterion

We study the ASHT problem from the perspective of minimizing the expected stopping time delay subject to constraints on the conditional probability of error, i.e., policies belonging to $\Delta(\alpha)$.

D. Lower Bound

The following proposition gives a lower bound for the mth moment of the conditional stopping time given hypothesis H_i for all policies belonging to $\Delta(\alpha)$.

Proposition 2.1: Assume (I). Then for any $\pi \in \Delta(\alpha)$ and any $m \geq 1$, we have

$$\inf_{\pi \in \Delta(\alpha)} E_i[\tau^m] \geq \left(\frac{|\log \sum_{j \neq i} \alpha_j|}{D_i} \right)^m (1 + o(1)) \; as \|\alpha\| \to 0. \tag{6}$$

where D_i is given in (5).

Proof: The proof for the case $m = 1$ follows from the proof of [2, Th. 2, p. 768] with minor modifications to account for the possibly different α_j. (See also proof of [8, Lem 2.1, Th 2.2]). For $m > 1$, the results follows from the fact that $E_i[\tau^m] \geq (E_i[\tau])^m$. ■

E. Achievability - Chernoff's Procedure A and related policies

Chernoff [2] proposed a policy termed *Procedure A* and showed that it has asymptotically optimal expected decision delay. Naghshvar and Javidi [12] proposed a class of policies (hereafter referred to as Policy-NJ), which are minor variants of *Procedure A*. We shall now argue that both the above policies are asymptotically optimal for all moments of stopping time. We consider only *Procedure A*. The result for Policy-NJ follows from a similar analysis.

Policy *Procedure A:* $\pi_{PA}(L)$
Fix $L > 0$.
At time n:

- Let $\theta(n) = \arg\max_i p_i(n)$
- If $p_{\theta(n)}(n) < \frac{L}{1+L}$, then $A_{n+1}(\mathbf{p})$ is chosen according to $\lambda_{\theta(n)}$, i.e., $P(A_{n+1}(\mathbf{p}) = a) = \lambda_{\theta(n)}(a)$.
- If $p_{\theta(n)}(n) \geq \frac{L}{1+L}$, then the test retires and declares $H_{\theta(n)}$ as the true hypothesis.

Policy $\pi_{NJ}(L)$:
Fix $0.5 < \tilde{p} < 1$. Fix $L > 0$.
At time n:

- If $0 \leq p_i(n) < \tilde{p}$ for every i, then $A_{n+1}(\mathbf{p})$ is chosen uniformly, i.e., $P(A_{n+1}(\mathbf{p}) = a) = \frac{1}{K}$.
- If $\tilde{p} \leq p_i(n) < \frac{L}{1+L}$ for some i, then $A_{n+1}(\mathbf{p})$ is chosen according to λ_i, i.e., $P(A_{n+1}(\mathbf{p}) = a) = \lambda_i(a)$.
- If $p_i(n) \geq \frac{L}{1+L}$ for some i, then the test retires and declares H_i as the true hypothesis.

As mentioned earlier, we focus only on *Procedure A*. However, we also consider the following variant of *Procedure A* for analysis. Policy $\pi^j_{PA}(L)$ is same as $\pi_{PA}(L)$ except that it stops only when $p_j(n) \geq \frac{L}{1+L}$.

Definition 2.2: We define the following entities:
1) $\tau^j_{PA}(L) := \inf\{n : p_j(n) \geq \frac{L}{1+L}\}$ is the stopping time at which the posterior belief of hypothesis H_j crosses the detection threshold for policy π^j_{PA}.
2) $\tau_{PA}(L) := \min_j \tau^j_{PA}(L)$ is the stopping time for policy $\pi_{PA}(L)$.
Observe that $\tau_{PA}(L) \leq \tau^j_{PA}(L)$ for all j.

Policy π_{PA} is defined to stop only when the posteriors suggest a reliable decision. This is formalized now.

Proposition 2.3: For Policy π_{PA}, the conditional probability of error under hypothesis H_i is upper bounded by

$$P_i(d \neq i) \leq \frac{1}{p_i(0)(1+L)}. \tag{7}$$

Consequently $\pi_{PA}(L) \in \Delta(\alpha)$ if for every i we have

$$\alpha_i \geq \frac{1}{p_i(0)(1+L)}. \tag{8}$$

We now proceed towards identifying the time-delay performance of the policy π_{PA}. Towards this, we first consider the easier to analyze policy π^i_{PA}. Note that as $L \to \infty$, we have $\alpha_i \to 0$, and therefore the policy has vanishing error cost. However, the time taken scales to infinity as given next.

Proposition 2.4: Assume (I) and (II). Consider policy π^i_{PA}. The following convergences then hold as $L \to \infty$:

$$\frac{\tau^i_{PA}(L)}{\log L} \to \frac{1}{D_i} \quad a.s. - P_i, \tag{9}$$

$$E_i\left[\left(\frac{\tau^i_{PA}(L)}{\log L} \right)^m \right] \to \frac{1}{D_i^m}. \tag{10}$$

A sandwiching argument on the likelihood ratio about the stopping time $\tau^i_{PA}(L)$ gives the first result. To show convergence in $\mathcal{L}_p$ we prove and use an exponentially boundedness property of the stopping time $\tau^i_{PA}(L)$. We omit the complete proof due to lack of space.

Theorem 2.5: Choose α as per (8). Assume (I) and (III). The family of policies $(\pi_{PA}(L))_{L>0}$ is asymptotically optimal in the sense that, for each $m \geq 1$, we have

$$\lim_{L\to\infty} \inf_{\pi \in \Delta(\alpha)} E^\pi_i \left[\left(\frac{\tau}{\log L} \right)^m \right] = \lim_{L\to\infty} E_i \left[\left(\frac{\tau_{PA}(L)}{\log L} \right)^m \right] = \frac{1}{D_i^m}. \tag{11}$$

Proof: By the choice of α and Proposition 2.1 we have

$$\lim_{L\to\infty} \inf_{\pi \in \Delta(\alpha)} E_i \left[\frac{\tau}{\log L} \right]^m \geq \frac{1}{D_i^m},$$

and by Proposition 2.4, we have the achievability of the above lower bound by policy $\pi_{PA}(L) \in \Delta(\alpha)$. ■

III. An Application to Visual Search

We now return to the visual search problem. In the visual search task, a subject has to identify an oddball image from amongst W images displayed on a screen (W = 6 in Figures 1(a) and 1(b)). With equal probability, the odd image can either be image I_1 or image I_2. When image I_1 is the odd image, the other $W-1$ images are of type I_2 and vice-versa. For the purpose of modeling, we make the following assumptions. The subject can focus attention on only one of the W positions, and that the field of view is restricted to that image alone. Further, we assume that time is slotted and of duration T. The subject can change the focus of his attention to any of the W image locations, but only at the slot boundaries. We assume that the subject would have indeed found the exact location of the odd image and the image identity before mapping it to a "left" or "right" decision. These are clearly oversimplifying assumptions, yet our model and analysis provide some insights into the visual search problem.

If the image in the focused location is I_i, we assume that a set of d neurons sensitive to these images produce spike trains, which constitute the observations. These are modeled as d Poisson point processes of duration T with rates $R_i = (R_i(1), R_i(2), \ldots, R_i(d))$.

The visual search problem can be formulated as a $2W$ hypothesis testing problem:

$$H_i, i \leq W: \text{ The odd image is } I_1 \text{ and is at location } i$$
$$H_i, i > W: \text{ The odd image is } I_2 \text{ and is at location } i - M$$

As the subject can change the point of focus at any slot boundary we have an ASHT problem.

We now calculate the optimal λ_i and the optimal neuronal metric D_i for the visual search problem. Recall that the set of controls $\mathcal{A} = \{1, 2, \ldots, W\}$. For notational simplicity let f_i denote the probability density on the observations when focusing on image I_i, $i = 1, 2$. The conditional probability density function q_i^a under hypothesis H_i when action a is chosen is :

$$q_i^a = \begin{cases} f_1 & i \leq W, a = i \\ f_2 & i \leq W, a \neq i \end{cases}$$
$$q_i^a = \begin{cases} f_2 & i > W, a = i - W \\ f_1 & i > W, a \neq i - W \end{cases}$$

Indeed, under Hypothesis H_i with $i \leq W$, the odd image is I_1 and is at location i. If the control is to focus on location i, i.e., $a = i$, then the subject views image I_1 and so the observations have density f_1 corresponding to I_1. The others are explained similarly.

The relative entropy between the probability densities for the various combinations of hypotheses and actions are:

(i) $\quad i \leq W, j \leq W$

$$D(q_i^a \| q_j^a) = \begin{cases} D(f_1 \| f_2) & a = i \\ D(f_2 \| f_1) & a = j \\ 0 & a \neq i, a \neq j, \end{cases}$$

(ii) $\quad i \leq W, j = i + W$

$$D(q_i^a \| q_j^a) = \begin{cases} D(f_1 \| f_2) & a = i \\ D(f_2 \| f_1) & a \neq i, \end{cases}$$

(iii) $\quad i \leq W, j > W, j \neq i + W$

$$D(q_i^a \| q_j^a) = \begin{cases} 0 & a = i \\ 0 & a = j - W \\ D(f_2 \| f_1) & a \neq i, a \neq j - W. \end{cases}$$

Our aim is to find the optimum λ_i that maximizes

$$D_i = \max_{\lambda} \min_{j \neq i} \sum_{a \in \mathcal{A}} \lambda(a) D(q_i^a \| q_j^a).$$

The corresponding optimum D_i thus obtained will yield the required "neuronal metric" (after a further scaling). Due to the symmetry in the problem, λ_i and D_i will be the same under all hypotheses $H_i, i \leq W$, and similarly it will be the same under all hypotheses $H_i, i > W$. Without loss of generality we consider the case $i \leq W$. Solution to the case when $i > W$ will have the same structure.

Theorem 3.1: Let $i \leq W$. The optimum λ_i and D_i are as follows

$$\lambda_i(i) = \frac{(W-3)D(f_2\|f_1)}{(W-1)D(f_1\|f_2) + (W-3)D(f_2\|f_1)},$$
$$\lambda_i(j) = \frac{D(f_1\|f_2)}{(W-1)D(f_1\|f_2) + (W-3)D(f_2\|f_1)} \qquad \forall j \neq i,$$
$$D_i = \frac{(W-2)D(f_1\|f_2)D(f_2\|f_1)}{(W-1)D(f_1\|f_2) + (W-3)D(f_2\|f_1)} \tag{12}$$

when $D(f_1\|f_2) > \frac{D(f_2\|f_1)}{W-1}$ and

$$\lambda_i(i) = 0, \lambda_i(j) = \frac{1}{(W-1)} \; \forall j \neq i, D_i = \frac{D(f_2\|f_1)}{(W-1)}$$

when $D(f_1\|f_2) \leq \frac{D(f_2\|f_1)}{W-1}$.

The proof is omitted due to lack of space. Note that the lower bound in Naghshvar and Javidi [12] would have $\max\{D(f_1\|f_2), D(f_2\|f_1)\}$ as an upper bound for the neuronal metric, but our bound tightens it by a factor of two.

d) Empirical Study on Neuronal Data: We now apply the results obtained in the previous section to the empirical data obtained from the experiments of Sripati and Olson [1]. Let T be the slot duration, during which one focuses attention on a particular image. $\mathcal{X}$ is the space of counting processes in $[0, T]$ with an associated σ-algebra. Let $P_{1,T}$ be the standard Poisson point process and let $P_{1,T}^{\otimes d}$ its d-fold product measure. Let $P_{\mathbf{R}_j,T}$ denote the probability measure P_j, so that $f_j = \frac{dP_{\mathbf{R}_j,T}}{dP_{1,T}^{\otimes d}}$. First we calculate the neuronal metric when f_1 and f_2 are vector Poisson processes of duration T with rates $\mathbf{R}_1 = (R_1(1), R_1(2), \ldots, R_1(d))$ and $\mathbf{R}_2 = (R_2(1), R_2(2), \ldots, R_2(d))$. Under an independence assumption, the relative entropy between the vector Poisson processes becomes the sum of relative entropies between the individual processes. Then

$$\begin{aligned} D\left(P_{\mathbf{R_1},T} \| P_{\mathbf{R_2},T}\right) &= E_{P_{\mathbf{R_1},\mathbf{T}}}\left[\log \frac{dP_{\mathbf{R_1},\mathbf{T}}}{dP_{\mathbf{R_2},\mathbf{T}}}\right] \\ &= \sum_{k=1}^{d} E_{P_{\mathbf{R_1(k)},\mathbf{T}}}\left[\log \frac{dP_{\mathbf{R_1(k)},\mathbf{T}}}{dP_{\mathbf{R_2(k)},\mathbf{T}}}\right] \\ &= T \sum_{k=1}^{d}\left[R_2(k) - R_1(k) + R_1(k) \log \frac{R_1(k)}{R_2(k)}\right]. \end{aligned}$$

Sripati and Olson [1] conducted the experiment with $W = 6$. The normalized per-unit-time per-neuron neuronal metric when image I_1 is the target image is

$$\tilde{D_1} = \frac{1}{dT} \frac{4D(P_{\mathbf{R_1},T} \| P_{\mathbf{R_2},T}) D(P_{\mathbf{R_2},T} \| P_{\mathbf{R_1}.T})}{5D(P_{\mathbf{R_1},T} \| P_{\mathbf{R_2},T}) + 3D(P_{\mathbf{R_2},T} \| P_{\mathbf{R_1},T)}}. \quad (13)$$

A similar relation holds when image I_2 is the target image.

The experimental data used in the empirical study consisted of the following. 1) Neuronal firing rate vectors were obtained from the IT cortex of rhesus macaque monkeys for twelve image pairs with the number of neurons ranging from 78 to 174 for different image pairs. 2) Reaction times statistics for detection of the odd image were obtained from experiments on human subjects. The neuronal behavioral index for an ordered pair of images (i, j) is the inverse of average decision delay minus a baseline delay. In Figure 2, we plot the behavioral discrimination index (speed of discrimination or inverse of time taken to decide) against the normalized per-unit-time per-neuron neuronal metric. The correlation between behavioral discrimination index and the neuronal metric was 0.90. This is smaller than the correlation of 0.95 between the behavioral discrimination index and the L_1 distance between the neuronal firing rate vectors. The discrepancy might arise from the highly simplified theoretical formulation, or from computational constraints in the brain. Nonetheless the close correspondence with the data suggests that the neuronal metric proposed here is a step in the right direction.

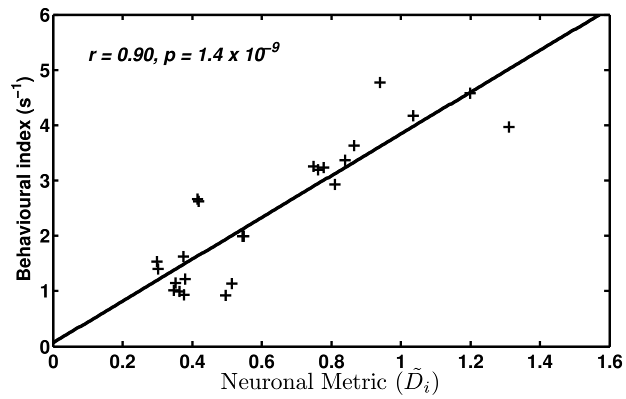

Fig. 2. The observed behavioral discrimination index is plotted against the neuronal metric calculated based on the observed neuronal responses. We found a striking degree of correlation $(r = 0.90, p = 1.4 \times 10^{-9})$.

Acknowledgments

This work is supported by the DST and UGC, Government of India. The authors would like to thank an anonymous reviewer for pointing our attention to [2] which greatly simplified our presentation.

References

[1] A. P. Sripati and C. R. Olson, "Global image dissimilarity in macaque inferotemporal cortex predicts human visual search efficiency," *Journal of Neuroscience*, vol. 30, no. 4, pp. 1258–1269, Jan 2010.

[2] H. Chernoff, "Sequential design of experiments," *The Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, 1959.

[3] A. Wald, "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, 1945.

[4] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *The Annals of Mathematical Statistics*, vol. 19, no. 3, pp. 326–339, 1948.

[5] A. Tartakovsky, "Sequential testing of many simple hypotheses with independednt observations," *Probl. Inform. Transm.*, vol. 24, no. 4, pp. 299–309, Oct 1988.

[6] C. Baum and V. Veeravalli, "A sequential procedure for multihypothesis testing," *IEEE Trans. Inf. Theory*, vol. 40, no. 6, pp. 1994 –2007, Nov 1994.

[7] V. Veeravalli and C. Baum, "Asymptotic efficiency of a sequential multihypothesis test," *IEEE Trans. Inf. Theory*, vol. 41, no. 6, pp. 1994 –1997, Nov 1995.

[8] A. Tartakovsky, "Asymptotic optimality of certain multihypothesis sequential tests: Non - i.i.d. case," *Statistical Inference for Stochastic Processes*, vol. 1, no. 3, pp. 265–295, 1998.

[9] V. Draglia, A. Tartakovsky, and V. Veeravalli, "Multihypothesis sequential probability ratio tests - part I. Asymptotic optimality," *IEEE Trans. Inf. Theory*, vol. 45, no. 7, pp. 2448 –2461, Nov 1999.

[10] M. Naghshvar and T. Javidi, "Active M-ary sequential hypothesis testing," in *IEEE Int. Symp. Info. Theory*, June 2010, pp. 1623 –1627.

[11] ——, "Information utility in active sequential hypothesis testing," in *48th Annual Allerton Conference*, Oct 2010, pp. 123 –129.

[12] ——, "Performance bounds for active sequential hypothesis testing," in *IEEE Int. Symp. Info. Theory*, Aug 2011, pp. 2666 –2670.

[13] R. P. N. Rao, "Decision making under uncertainty: A neural model based on partially observable markov decision processes," *Frontiers in Computational Neuroscience*, vol. 4, no. 00146, 2010.