

Stopping Criterion Design for Recursive Bayesian Classification: Analysis and Decision Geometry

Aziz Koçanaoğulları, Murat Akcakaya, Deniz Erdoğan

Abstract—Systems that are based on recursive Bayesian updates for classification limit the cost of evidence collection through certain stopping/termination criteria and accordingly enforce decision making. Conventionally, two termination criteria based on pre-defined thresholds over (i) the maximum of the state posterior distribution; and (ii) the state posterior uncertainty are commonly used. In this paper, we propose a geometric interpretation over the state posterior progression and accordingly we provide a point-by-point analysis over the disadvantages of using such conventional termination criteria. For example, through the proposed geometric interpretation we show that confidence thresholds defined over maximum of the state posteriors suffer from stiffness that results in unnecessary evidence collection whereas uncertainty based thresholding methods are fragile to number of categories and terminate prematurely if some state candidates are already discovered to be unfavorable. Moreover, both types of termination methods neglect the evolution of posterior updates. We then propose a new stopping/termination criterion with a geometrical insight to overcome the limitations of these conventional methods and provide a comparison in terms of decision accuracy and speed. We validate our claims using simulations and using real experimental data obtained through a brain computer interfaced typing system.

Index Terms—Active Learning, Sequential Decision Making, Recursive Bayesian Classification, Optimal Stopping Criterion Design

1 INTRODUCTION

Recursive Bayesian inference for classification (RBC) is beneficial in gradually increasing decision quality by incorporating more evidence into the decision process in scenarios where data or evidence is acquired sequentially over time. The most recent belief, represented by the latest label posterior probability distribution, is obtained by incorporating new evidence in a Bayesian manner at each update step [1], [2], [3], [4]. The trade-off between decision confidence and evidence acquisition cost is controlled by a stopping criterion that controls when to terminate evidence collection and return an estimated class label. Fundamental components of RBC include: (*S*) A stopping criterion based on the posterior probability to stop evidence collection; (*Q*) a querying step to decide how to collect further evidence from relevant sources to benefit speed and accuracy objectives of RBC; (*C*) a classification objective based on the posterior distribution and loss values attributed to each true label and decision option pair to determine the optimal decision once the stopping criterion has been satisfied. The iterative process can be summarized as follows:

while (*S*) not satisfied { **do** (*Q*) }; **return** state with (*C*)

- Aziz Koçanaoğulları and Deniz Erdoğan were with Northeastern University Department of Electrical and Computer Engineering 409 Dana Research Center 360 Huntington Avenue Boston, MA 02115.
E-mail: akocanaoğullari@ece.neu.edu, erdogmus@ece.neu.edu
- Murat Akcakaya was with Pittsburg University Department of Electrical and Computer Engineering 1238 Benedum Hall Pittsburgh, PA 15261.
E-mail: akcakaya@pitt.edu

This work is supported by NSF (IIS-1149570, CNS-1544895, IIS-1715858, IIS-1717654, IIS-1844885, IIS-1915083), DHHS (90RE5017-02-01), and NIH (R01DC009834).

In this paper, we focus on designing stopping criteria for (*S*) based on the latest label posterior distribution $p = [p_1, p_2, \dots, p_n]$, which is a categorical probability distribution. In the remainder of this paper, in order to keep the illustrations and derivations simple, we assume that the objective in (*C*) is to minimize probability of error (therefore 0-1 loss is assumed), thus when a decision is made, it will be based on the maximum a posteriori (MAP) classification rule; that is the then-most-likely-label in the latest posterior will be selected as the decision. The presented approach can be generalized to the more general expected loss minimization classification setting [5], [6], but is outside the scope of this paper. We also do not discuss different querying strategies [7], [8], [9], [10], as they too are outside the scope.

The RBC procedure will terminate when the stopping criterion in (*S*) (namely ($\mathcal{S}_C$)) is met by the current label posterior distribution ($\mathcal{S}_C(p) = \text{true}$). The most commonly used approach is to require that a confidence threshold has been exceeded ($\mathcal{S}_C(p) = \text{true}$ if $\max_i p_i > \tau$). All label probability distributions p that satisfy $\mathcal{S}_C$ form a set called the stopping region $\mathcal{S}_R = \{p | \mathcal{S}_C(p) = \text{true}\}$. If the posterior distribution falls into this set, the system terminates and a classification-decision is made based on (*C*). Uncertainty termination rule is another canonical stopping decision that utilizes classification uncertainty (e.g. $\mathcal{S}_C(p) = \text{true}$ if $H_\alpha(p) < c_{H_\alpha}$). Disfavoring candidate categories (assigning small probability masses in the latest posterior) will yield lower classification uncertainty. However, this is independent of the classification confidence level and hence can sacrifice accuracy. In this paper, we analyze $\mathcal{S}_C$ with the aid of analytical representations that arise from the geometry of the probability simplex. A probability simplex for three categories ($\{a, b, c\}$) is illus-

Stopping Region: $S_R := \{p \mathcal{S}_C(p) = \text{true}\}$			
Identifier	Root	Stopping Criterion ($\mathcal{S}_C$)	Reference
(M1)	Confidence Level	$\max_{\sigma} p(\sigma \mathcal{H}_s) > \tau$	[11]
(M2)	Renyi Entropy ($\alpha = 2$)	$H_{\alpha}(p(\sigma \mathcal{H}_s)) < c_{H_2}$	[12]
(M3)	Shannon's Entropy	$H(p(\sigma \mathcal{H}_s)) < c_H$	[13]
(M4)	Renyi Entropy ($\alpha = 0.2$)	$H_{\alpha}(p(\sigma \mathcal{H}_s)) < c_{H_{0.2}}$	[12]
(M5)	Kullback Liebler Div.	$\delta_{KL}(p(\sigma \mathcal{H}_s), p(\sigma \mathcal{H}_{s-1})) < c_{\delta_{KL}}$	[14]

TABLE 1: Conventional stopping criteria, here $c_s \in \mathbb{R}$ denote the limits of the criteria. These methods are further used in the experiments section with the respective identifiers as a baseline to the proposed perspective of the stopping method. Limiting criterion $\mathcal{S}_C$ defines a set with all distributions p that satisfy the condition. In recursive estimation, if the posterior probability becomes an element of the set, the system terminates to make an inference/classification.

trated in Fig. 2. Confidence and entropy decision boundaries are marked in the illustration with (M1) and (M3) respectively. Especially, to achieve high accuracy in classification, a high confidence threshold is usually preferred. However, based on the prior information if the true class is one of the least favored, confidence on true class increases gradually and hence through RBC the system sacrifices speed. This is illustrated in Fig. 2-(a) where the prior probability starts away from the corner of interest and gradually gets closer to a . Uncertainty based approaches are more flexible in termination but sensitive to number of classes. Considering drawbacks of conventional stopping criteria, it is preferred to optimize stopping in different RBC scenarios. The design of an optimal stopping region (S_R , see Table 1) has been referred to as "optimal stopping criterion design" [15]. Optimum S_R design has many practical applications. For example, in any resource constrained state estimation or recursive classification setting, stopping criteria could be preferred to be optimized to minimize time requirements of the system with a marginal expense in estimation confidence/accuracy. Some specific example applications include but not limited to, radar tracking with potential adaptive waveform design and radar subset selection or optimum radar placement [16], [17], [18], brain-computer interfaced typing systems [19], text recognition from video stream [20].

In this paper we propose a stopping criterion building on the strengths of the existing conventional criteria but improving their shortcomings. That is, the method we propose is robust to potential decrease in accuracy while promoting early stopping. Motivated by problem geometry and RBC mechanics, we propose a stopping criteria function $\mathcal{S}_C(p)$ that yields a bent-in stopping boundary towards the uniform distribution (u_n in Fig.2-(d)) as illustrated by the blue curves in Fig.2-(c). Unlike conventional methods, this enables early stopping decisions especially in the absence of a strong contender. When there exists a strong contender (posterior moving near the edges of simplex) it does not sacrifice accuracy.

Problem Formulation and Related Work

Let class label σ be an element of a finite set $\mathcal{A}$. We refer to each iteration in RBC as a *sequence*, indexed by $s \in \mathbb{N}$. Each sequence may include a set of measurements acquired in response to queries $\Phi_s \triangleq \{\phi_s^1, \dots, \phi_s^N\}$, where N denotes number of queries. These measurements provide evidence (raw data or processed features) $\epsilon_s \triangleq \{\epsilon_s^1, \dots, \epsilon_s^N\}$ conditioned on the queries, and the true label. RBC posterior updates use this evidence. For notation simplicity, we use $\mathcal{H}_s \triangleq \{\epsilon_{1:s}, \Phi_{1:s}, \mathcal{H}_0\}$ to represent the combination of all evidence collected in sequences 1 to s , as well as the prior, $\mathcal{H}_0$.

The relationship between RBC variables is illustrated in Fig. 1. Through this paper we assume evidences conditioned on state and query tuples are independent of each other $p(\epsilon_i, \epsilon_j | \sigma, \Phi_i, \Phi_j) = p(\epsilon_i | \sigma, \Phi_i) p(\epsilon_j | \sigma, \Phi_j) \forall i \neq j$. There are no limiting assumptions that are set on prior information $p(\sigma | \mathcal{H}_0)$. The importance of the prior information in the RBC paradigm is that as prior probability is the initial confidence assessment of the estimation, it determines the evolution of the estimation as it will be discussed in Sec.3.2. RBC steps (assuming MAP classification rule) are presented in the following:




Fig. 1: Proposed probabilistic graphical model representing the s^{th} sequence for a RBC.

$$\begin{aligned}
 (S) : & \quad p(\sigma | \mathcal{H}_s) \in (S_R := \{p | \mathcal{S}_C(p) = \text{true}\}) \\
 (C) : & \quad \hat{\sigma} = \arg \max_{\sigma \in \mathcal{A}} p(\sigma | \mathcal{H}_s) \\
 (Q) : & \quad \Phi_{s+1} \rightarrow \epsilon_{s+1} \text{ with the anticipated joint} \\
 & \quad p(\sigma | \mathcal{H}_{s+1}) = p(\sigma | \mathcal{H}_s) \frac{p(\epsilon_{s+1} | \sigma, \Phi_{s+1})}{p(\epsilon_s | \Phi_{s+1})}
 \end{aligned} \tag{1}$$

In this paper we focus on (S). Specifically, we discuss the limitations of the conventional criteria summarized in Table 1, and propose a new family of stopping criteria, then illustrate theoretically, and with numerical examples and experiments, the benefits of the proposed approach. The most common $\mathcal{S}_C$ is thresholding the highest value in the latest posterior by τ : $\mathcal{S}_C(p) = \text{true}$ if $\max_i p_i > \tau$ (M1) [11], [21]. This rule directly enforces that a decision is made with a pre-specified confidence level, and does not consider the distribution of probability mass among other options. To minimize false alarm rate, the threshold is usually set to a large number. This prevents early stopping; and to achieve high accuracy, furthermore leads to redundant evidence collection [22]. On the other hand, information theoretic objectives utilize more than a single category in termination decision. Unlike confidence based approaches, uncertainty based stopping allows early stopping if the system uncertainty is below a desired threshold where the confidence is not yet high enough for confidence stopping. Uncertainty measures, including Shannon entropy [23], can be used to have termination based on the spread of posterior probability mass. Golovin [14] uses Shannon entropy as a stopping criterion: " $\mathcal{S}_C(p) = \text{true}$ if $H(p) < c_H$ " (M3). Yingzhen [24] and Igal [12] use Renyi entropy: " $\mathcal{S}_C(p) = \text{true}$ if $H_{\alpha}(p) < c_{H_{\alpha}}$ ". Uncertainty based stopping criteria based on Renyi entropy $H_{\alpha}(p)$ family with $\alpha \geq 0$ includes, as special cases, Shannon entropy (in the limit as $\alpha \rightarrow 1$), and confidence thresholding (in the limit as $\alpha \rightarrow \infty$) methods




Fig. 2: Decision geometry for recursive classification where true class is a . (a) The initialization of the system with a prior probability (p^0) and (S) region. The system terminates once the posterior yields within the (S) region (dashed). (b) Stopping boundaries for $(M1)$ and $(M3)$ are visualized. Highlighted region is the decision region for ' a ' enforced by (C) . (c) Three different perspectives in S_R design (entropy:diagonal-stripes/top, confidence:vertical-stripes/left, proposed:blue-stripes/right), we want to tilt the confidence boundary towards u_n as illustrated with dashed arrows to achieve early stopping. (d) The origin of the simplex u_n and special points defined in (4). Equi-entropy contours and corresponding posterior threshold lines are in solid black. The intersection point is at v_n (See Observation 1).

$(M2)$, $(M4)$. It is possible to terminate if the changes in posterior distribution between sequences become smaller over iterations, Weinshall [25] and Geisser [13] apply a threshold to the Kullback-Leibler (KL) divergence between two consecutive posterior distributions to terminate evidence collection: " $\mathcal{S}_C(p) = \text{true}$ if $D_{KL}(p_s) || p_{s-1} < c_{D_{KL}}$ ", therefore terminate RBC when the distance between two consecutive classification posteriors becomes small $(M5)$. Pavlichin [26] proposes a chained-KL divergence to monitor posterior progression in the probability simplex. Banerjee [27] proposes using Bregman divergences in the context of clustering. All of these stopping criteria focus on how far the posterior is from the uniform distribution that is at the center of the probability simplex, as opposed to assessing how close the posterior to a vertex of the simplex (corresponding to a one-hot distribution, see Fig. 2-(b)) [28].

Contributions

We introduce a geometrical representation for recursive Bayesian classification in Sec. 2. In Sec. 3 we show that: (i) uncertainty based methods are sensitive to the number of possible classes, and (ii) conventional approaches ignore the posterior update trajectory in the recursive classification task. Using (i) and (ii) together, we show that the stopping regions defined based on uncertainty methods diverge from the region formed by confidence threshold defined over the posterior distribution resulting in decrease in classification accuracy. In Sec. 4, we propose a stopping criterion approach which utilizes the strengths of conventional stopping criteria but robust to their drawbacks. Our approach is also inline with the progression of the posterior update (blue curves in Fig.2-(c)) to achieve this balance between speed and accuracy compared to other approaches. Throughout Sec. 5 first with synthetic experiments we show that our approach allows RBC to stop earlier where there is no close competitor. This results in faster decision cycles without sacrificing accuracy. Therefore, our approach maintains a high classification accuracy where uncertainty based methods stop early with a significant accuracy loss. We also include results from a brain-computer interfaced typing system in which the user is trying to type a letter with electroencephalogram (EEG) responses to stimuli. Due to observation noise, multiple EEG acquisitions are required. Accurate Bayesian classification

of the target letter relies on prior information provided by a language model and recursive evidence collection. The language model and the EEG characteristics of the user determines the path to classification in simplex. We show that with the proposed method, it is possible to commit to a letter sooner with marginal accuracy loss.

Detailed demonstrations and proofs of the analytical results are provided in the appendix Sec. 7.4 for neat presentation.

2 PROBLEM GEOMETRY

To promote the representative power of visualization, in this paper we use information geometric representation of recursive classification problem [29]. As will be discussed, this representation allows us to represent progression of the posterior probability distribution $p(\sigma | \mathcal{H}_s)$ as the evidence collection steps s increases and to introduce a mean for analytical reasoning for stopping criterion $\mathcal{S}_C$ design. We use simplex as a domain for probability distributions;

$$\Delta_n = \{(p_1, p_2, \dots, p_n) \in \mathbb{R}^n | p_i > 0 \forall i, \sum_i p_i = 1\} \quad (2)$$

Here Δ_n represents the set that includes categorical probability distributions, where the simplex is an $n - 1$ dimensional geometrical object in $\mathbb{R}^n$. We visualize an example simplex for a category with 3 elements in Fig. 2. To preserve neat visualization, we interpret our reasoning using the triangle throughout the paper, but all the results can be generalized to n -dimensional case without loss of generality. Moreover, we use the following addition operation that allows us to represent posterior updates;

Addition: Given $p, q \in \Delta_n$ the addition operation is defined as:

$$p \oplus q = \frac{[p_1 q_1, p_2 q_2, \dots, p_n q_n]}{\sum_i p_i q_i} \quad (3)$$

Observe that addition satisfies the Bayes' theorem: $p(\sigma | \varepsilon) = p(\sigma) \oplus p(\varepsilon | \sigma)$ which allows us to represent the evolution of the posterior distribution using the addition operation within the simplex. For algebra to work $p(\varepsilon | \sigma)$ does not need to be unit ℓ_1 norm and hence $\notin \Delta_n$, but what we interpret here is the normalized vector over σ even algebraically $\oplus$ is applicable. Rigorous definition of this simplification is shown

in the appendix Sec. 7.1. The zero element in Δ_n is denoted by u_n being the uniform distribution as shown in Fig. 2-(d).

Proposition 1 ([30]). *Simplex $\{\Delta_n | \oplus, \otimes\}$ forms a vector space.*

In the presence of a prior distribution, recursive classification starts with a distribution probably different than uniform distribution as presented in Fig. 2-(a). To make a correct decision, the posterior for respective element is required to be the most likely (e.g. the region for a is highlighted in Fig. 2-(b)). Through evidence collections, the posterior distribution moves within the simplex visualized in Fig. 2-(c) that is algebraically denoted with $\oplus$ above. (S) criterion forms the decision lines within the simplex and the inequality condition covers the area of evidence collection termination that are visualized by dashed areas in Fig. 2. Once the posterior probability distribution reaches that area, the system terminates evidence collection, makes an inference outputs the classification decision.

3 CONVENTIONAL STOPPING CRITERIA

In this section we discuss the limitations of the conventional methods that are used for stopping criteria $\mathcal{S}_C$ design. Consider the following 2 motivating examples assuming 10 class classification problems;

example1: $p, q \in \Delta_{10}$ (i.e. 10 dimensional probability mass functions) such that $p = [.6, .4 - 8\epsilon, \epsilon, \dots]$ where $0 < \epsilon \ll 1$ and $q = [.7, .03, \dots]$. It is apparent that q is a better stopping point to decide on the 1st element as the classification result. However, if Shannon's entropy is used to measure the distance from uniform distribution u_n , and accordingly utilized as a stopping criterion, then $0.97 \approx H(p) < H(q) \approx 1.82$. This means that the evidence collection would have been terminated when p is reached. Hence, the uncertainty based (S) suffers from not taking the probability mass index of interest into consideration.

example2: $p, q, x, y \in \Delta_{10}$ with $p = [.5, .5 - 8\epsilon, \epsilon, \dots]$, $q = [.6, .4 - 8\epsilon, \epsilon, \dots]$ where $0 < \epsilon \ll 1$ and $x = [.5, 0.05, \dots]$, $y = [.6, 0.04, \dots]$ and we compare p with q then x with y . Assume that p evolved into q and x evolved into y through Bayesian recursions. Therefore, the propagation from p to q , and x to y , correspond to different paths on the probability simplex. Specifically, from p to q , the path is close to the edges (See Fig. 2-(b) for 3 class example with $[0.5, 0, 0.5]$), and the path from x to y is close to the center. Observe that the confidence level for the true class (the 1st class) increased the same amount (.1) for both of these cases. However for p and q the 2nd best class is still a close competitor having almost all the remaining probability mass. Whereas for x and y there is no close competitor. Therefore, even for these two cases the confidence levels for the true class is the same, one can imply that the in case 2 (x, y) we are still confident there exists no strong competitor. This means that the location of posterior distributions in Δ_n matters and with an appropriate design of stopping criterion, the system evidence collection could be stopped when there exists no close competitor.

In this section we analytically show what these examples mean for the stopping criterion ($\mathcal{S}_C$) design.




Fig. 3: Effect of number of possible categories (different color/line codes) on the decision boundary. (a) represents the difference between entropy values of the probabilities v_n, w_n for a given confidence level τ . (b) represents the τ and $\hat{\tau}$ values where $H(w_n(\hat{\tau})) = H(v_n(\tau))$.

3.1 Frail Confidence

The confidence behavior of uncertainty based methods are similar. Shannon entropy on the other hand is the most commonly used and hence in this section we specifically focus on (M3), as defined in Table.1. We refer the reader to the appendix Sec. 7.2 for decision boundary similarities between (M2)-(M3) and (M4). To point out the relationship between the confidence in classification and entropy based stopping region $\mathcal{S}_R(H(\cdot)) := \{p | H(p) < c_H\}$ (M3) we define two special probability points in Δ_n ;

$$v_n(\tau) = \left[\tau, \frac{1-\tau}{n-1}, \dots, \frac{1-\tau}{n-1} \right] \quad (4)$$

$$w_n(\tau) = [\tau, 1-\tau, 0, \dots, 0]$$

Here, $v_n(\tau)$ corresponds to distributions where one class has likelihood of τ and the others share the remaining probability uniformly, $w_n(\tau)$ corresponds to distributions where only two class exist with probability values $\tau, 1-\tau$. These definitions are also visualized in Fig. 2-(d).

In this paper, we bundle the uncertainty and confidence τ via $\tau' = H(v_n(\tau))$. We show the following weak points: $H(\cdot)$ is sensitive to n and τ , and $H(\cdot)$ conflicts with the inference step (C) of Eqn. (1) for a given set of parameters. To analyze these, we introduce the relationship between confidence and entropy $\mathcal{S}_C$ s with the help of (4). We state for a confidence τ that entropy achieves its maximum value at the point $v_n(\tau)$ and for an equi-entropy contour with value of $H(v_n(\tau))$ maximum achievable confidence is τ . See the following proposition.

Proposition 2. *For a defined confidence level τ ;*

$$C_\tau = \{p | \max_i p_i = \tau\} \implies \max_{p \in C_\tau} H(p) = H(v_n(\tau))$$

$$S_\tau = \{p | H(p) = H(v_n(\tau))\} \implies \max_{p \in S_\tau} \max_i p_i = \tau \quad (5)$$

Observation 1. Following descriptions in Proposition 2, confidence line for τ , C_τ , intersects with the equi-entropy contours, S_τ , only at $v_n(\tau)$ points.

Following this observation, we state that the set between the corners of the simplex Δ_n and C_τ is a subset of the set between the corners and S_τ as presented in the following observation:

Observation 2. Define $S_1 = \{p | \max_i p_i \geq \tau\}$ and $S_2 = \{p | H(p) \leq \tau'\}$, $\tau' = H(v_n(\tau))$, $\forall \tau \in [1/n, 1]$ then $S_1 \subset S_2$.

Therefore the (S) region designed by entropy S_1 is larger than the region designed with the confidence threshold S_2 .

It is apparent that the stopping region defined by entropy (M3) is larger than the stopping region defined through confidence level (M1) and hence enlarging the $\mathcal{R}_S$ in (1). This increase in the region, on the other hand, decreases confidence. To analyze it we need to find the minimum confidence in an equi-entropy contour. It is shown that the minimum confidence is attained at $w_n(\tilde{\tau})$ that satisfies the following;

Observation 3. [31]

$$\begin{aligned} H(v_n(\tau)) < 1, S_\tau = \{p|H(p) = H(v_n(\tau))\} \\ \text{then } \min_{p \in S_\tau} \max_i p_i = \tilde{\tau} \\ \text{s.t. } -\tilde{\tau} \log_2(\tilde{\tau}) - (1 - \tilde{\tau}) \log_2(1 - \tilde{\tau}) = H(v_n(\tau)) \end{aligned} \quad (6)$$

This observation states that the entropy based (S) boundary attains the minimum required max-probability value $\tilde{\tau}$ at $w_n(\tilde{\tau})$. $\tilde{\tau}$ vs τ and entropy values difference between $w_n(\tau)$ and $w_n(\tilde{\tau})$ are presented in Fig. 3 for changing number of categories in the classification. Observe from Fig. 3 that entropy is fragile with respect to the number of classes and the difference between $\tilde{\tau}$ and τ increases by decreasing values of τ . Also observe that $H(w_n(\tilde{\tau})) - H(w_n(\tau))$ decreases linearly with τ , and increases exponentially (dotted-line) with n as presented in Fig. 3-(a).

By nature, uncertainty based (S) is capable of returning a class label when the confidence of the class is low as shown in Observation 3. If there exists no $\tilde{\tau}$ that satisfies Observation 3, equi-entropy contours do not intersect with the simplex boundary. This condition might result in immediate stopping if one of the classes is already unfavorable as shown in Fig. 4-(a) with the example trajectories close to $[a, c]$ edge.

The above Proposition 2, and Observations 1-3 show that under certain conditions, the regions defined by uncertainty criteria may significantly diverge from the region formed by the confidence level threshold that was defined over the posterior distribution. To avoid such drawbacks, the system is usually set to a high confidence level (e.g. $\tau \approx .95$), such termination already mandates system continue with redundant recursions of evidence collection to achieve high confidence. Accordingly, despite their differences, uncertainty based stopping (M3) and constant confidence threshold (M1) yield similar performances in recursive classification tasks. To analyze these similarities in the classification performances, in the next section, we investigate the behavior of the posterior motion over the probability simplex. Through such an analysis, we then gain an insight into designing stopping criterion that will avoid high confidence levels and corresponding redundant recursions/evidence collections.

3.2 Posterior Motion

Given a prior point $p(\sigma)$, the posterior evolution with given pairs of queries and evidences (Φ, ε) ,

$$p(\sigma|\Phi_{0:s}, \varepsilon_{0:s}) = p(\sigma) \oplus p(\varepsilon_1|\sigma, \phi_1) \oplus \dots \oplus p(\varepsilon_s|\sigma, \phi_s)$$




Fig. 4: Monte Carlo simulated trajectories for probability evolution in 3D simplex. To simulate the trajectories we sample evidence from lognormal distributions. (a) Each color represent 100 simulated examples from a different starting point and dashed lines representing the means of the trajectories. Observe the case where one of the classes is already disfavored (one between a, c), if decision boundary does not intersect with the edge, that results in immediate termination. (b) Three different perspectives for S_R design. We observe that the behavior of entropy (black) and confidence threshold (blue) are similar if trajectories are considered. In the method section we propose a region design using an equi-distance curves wrt. to the corners (red).

Posterior trajectory is the list of vectors that is obtained within RBC over sequences of evidence collections: $[p(\sigma), p(\sigma|\Phi_0, \varepsilon_0), \dots, p(\sigma|\Phi_{0:s}, \varepsilon_{0:s})]$. Posterior motion refers to the changes in RBC over sequences that resulted in the posterior trajectory. In this section we aim to show that the trajectory of the posterior distribution follows a central path (Proposition 3) as it gets closer to the corner of interest as illustrated in Fig. 4. This will further allow us to clarify the similar early stopping behavior of (M1) and (M3) geometrically. By understanding motion we will design a stopping condition that promotes early stopping while preventing accuracy loss for a motion on edges.

The motion is determined by the evidence likelihood $p(\varepsilon_s|\sigma, \phi_s)$ at each step s and the evidence collected through query ϕ_s . In this section we assume two noisy information channels for true class and incorrect class respectively. Once the true class is queried, evidence is sampled from the "positive" distribution and evidence is sampled from "negative" distribution otherwise and we assume the posterior probability of the true state/class increases on average [10].

Lemma 1. Let $a \in \mathcal{A}$ and let $\phi(a)$ denote the query related with state a then $p(\sigma) = [p_1 = p(a), p_2, \dots, p_n]$. $p(\sigma)$, $[1, 0, \dots, 0]$ and $p(\sigma) \oplus p(\varepsilon|\sigma, \phi(a)) \forall \varepsilon, \forall p(\sigma)$ are collinear.

Hence, once the system queries the environment, posterior probability for classification takes a step on the line that passes through the current position and the corner addressed by query. Moreover, on average, if the query addresses true state, the posterior moves towards the respective corner, if the query addresses an incorrect estimate the posterior moves away from the respective corner. To show that in general posterior gets closer to a central position, we need the projection of a posterior point to the line that passes through the center (u_n) and one of the corners.

Lemma 2. Given $p(\sigma) = p = [p_1, p_2, \dots, p_n] \in \Delta_n$ and given the line $\ell_{c_n, i=1} = \{[\tau, \frac{1-\tau}{n-1}, \dots, \frac{1-\tau}{n-1}] | \tau \in [0, 1]\}$ then ℓ_2

norm projection of point $p(\sigma)$ onto line $\ell_{c_n, i=1}$ is the following:

$$\begin{aligned} \text{proj}_{\ell_{c_n, i=1}}(p(\sigma)) &= \arg \min_{p \in \ell_{c_n, i=1}} \|p(\sigma) - p\|_2 \\ &= \left[p_1, \frac{1-p_1}{n-1}, \dots, \frac{1-p_1}{n-1} \right] \end{aligned} \quad (7)$$

To give an example, line $\ell_{c_n=3, i=1}$ is visualized in Fig. 2-(d) where 1st location is a . With this projection operation we show that ℓ_2 -norm distance between the projection and the actual point d_s at sequence s decreases quadratically with respect to the posterior as s increases. This statement is given in the following proposition:

Proposition 3. *Following Lemma 2, given $p(\sigma) \in \Delta_n$:*

$$\|p(\sigma) - \text{proj}_{\ell_{c_n, i=1}} p(\sigma)\|_2^2 \propto (1-p_1)^2$$

Following this, one defines the reduction between two sequences,

$$\begin{aligned} \exists \hat{s} \text{ s.t. } \|p(\sigma|\mathcal{H}_s) - \text{proj}_{\ell_{c_n, i=1}} p(\sigma|\mathcal{H}_s)\|_2 &= d_s \\ \text{where } \tilde{p}(\sigma|\mathcal{H}_s + 1) &= p(\sigma|\mathcal{H}_s) \oplus E_\varepsilon(p(\varepsilon|\sigma, \phi)) \quad \forall \phi \\ \|\tilde{p}(\sigma|\mathcal{H}_{s+1}) - \text{proj}_{\ell_{c_n, i=1}} \tilde{p}(\sigma|\mathcal{H}_{s+1})\|_2 &= d_{s+1} \\ \text{Using Lemma 1} \implies d_s &> d_{s+1} \quad \forall s \geq \hat{s} \end{aligned}$$

In recursive classification, on average, probability of the true state/class increases sequentially and following the proposition, the posterior probability gets closer to the line that passes through the respective corner and the center. We visualize examples of average trajectories in Fig. 4-(a). Additionally we presented in Fig. 4-(b) that unlike (M1) $\mathcal{S}_C$ curves (blue), (M3)-curves (black) allow additional $\mathcal{S}_R$ (shaded areas) that is concentrated closer to edges and they promote early stopping. Considering both Fig. 4-(a,b) together, if posterior follows a central path (M3) behaves as (M1) and (M3) cannot promote early stopping. Next, we propose a new stopping criterion perspective to extend the stopping region centrally towards the uniform distribution (e.g. red boundary in Fig. 4-(b)). This approach will address our observations in Sec. 3.1 and Sec. 3.2 that the new approach will be robust to number of categories while promoting early stopping with marginal accuracy loss.

4 PROPOSED PERSPECTIVE

In the previous sections we argued equi-entropy contours formed for (M3) are centered around u_n . Moreover (M3) is sensitive to number of categories and stagger in cases where some of the classes are already unfavorable. Additionally we propose to bend (M1) from the center towards u_n e.g. the red boundary in Fig. 4-(b). Trivially, it is possible to form this by equi-distant points to the respective corners. However, by definition, edges and corners are $\notin \Delta_n$ (represents ∞ [30]) which prevents measuring the distance with conventional information theoretic approaches. To avoid this, exploiting $\Delta_n \subset \mathbb{R}^n$, one can use a distance measure δ defined over $\mathbb{R}^n$ (e.g. ℓ_p norms) and intersect the $\bar{\tau}$ ball around p $\tilde{B}_\delta^\tau(p) = \{x|\delta(p, x) < \bar{\tau} \text{ with } \Delta_n \text{ to obtain } \tilde{B} = B \cap \Delta_n \text{ centered around a corner and } \in \Delta_n. \tilde{B}_\delta^\tau(p) = \{x|\delta(p, x) < \bar{\tau}, x \in \Delta_n\}$. Let $c^k \in \Delta_n$ be the k^{th} corner (e.g. $c^1 = [1, 0, \dots, 0]$), decision region with δ and $\bar{\tau}$ is defined as;

$$\mathcal{S}_R : p \in \bigcup_{k \in \{1, 2, \dots, n\}} \tilde{B}_\delta^\tau(c^k) \quad (8)$$

We use a distance measure influenced by Kittler's work [32]: the *delta* divergence. We use the definition of *delta*-divergence and modify it to define our novel decision region. This divergence determines an interest set (e.g. indices of the most probable elements in respective distributions) and compare the probability mass in these elements. We limit the set to two elements in our paper. The corresponding distance between $p, q \in \Delta_n$ is denoted as δ_{MP} (MP indicating 'method proposed') defined as the following;

$$\begin{aligned} \delta_{MP}(p, q) &= \sum_{i \in I} |p_i - q_i|, \quad I = \{j_1, j_2, k_1, k_2\} \\ j_1 &= \arg \max_i p_i, j_2 = \arg \max_{i \neq j_1} p_i \\ k_1 &= \arg \max_i q_i, k_2 = \arg \max_{i \neq k_1} q_i \end{aligned} \quad (9)$$

Observe the measure satisfies non-negativity, identity of indiscernibles and symmetry. Using this distance to obtain balls, we can state the following proposition about the stopping region and criterion, $\mathcal{S}_R$ and $\mathcal{S}_C$ respectively;

Proposition 4 (Proposed $\mathcal{S}_R$ (MP)). *c^k being k^{th} corner, $p \in \Delta_n, \delta = \delta_{MP}, \bar{\tau} \in [1/n, 1]$ then;*

$$\begin{aligned} \mathcal{S}_R : p \in \bigcup_{k \in \{1, 2, \dots, n\}} \tilde{B}_\delta^\tau(c^k) \\ \equiv \mathcal{S}_C : p_{j_1} - p_{j_2} > 1 - \bar{\tau} \end{aligned} \quad (10)$$

where $j_1 = \arg \max_i p_i$ and $j_2 = \arg \max_{i \neq j_1} p_i$

We denote this method by (MP). We visualize an example decision boundary in Fig. 4-(b). As can be observed from the Fig. 4-(b), respective decision boundary is inline with the motion of the probability distribution. To have a decision boundary bending the boundary at τ for (M1) from its center the following condition is required;

Observation 4. Given $\bar{\tau} = 2 - 2\tau$ and WLOG for $p \in \Delta_n$, $\arg \max_i p_i = 1$. Define (MP) and (M1) decision boundaries;

$$\begin{aligned} C_\tau &= \{p|p_1 = \tau\} \\ B_{\bar{\tau}} &= \{p|p_1 - p_m = 1 - \bar{\tau}, m = \arg \max_{i \neq 1} p_i\} \end{aligned} \quad (11)$$

$$\implies C_\tau \cap B_{\bar{\tau}} = w_n(\tau) \text{ and } \max_{p \in B_{\bar{\tau}}} \max_i p_i = \tau$$

Observation 5.

$$\begin{aligned} B_{\bar{\tau}} &= \{p|p_1 - p_m = 1 - \bar{\tau}, m = \arg \max_{i \neq 1} p_i\} \\ \min_{p \in B_{\bar{\tau}}} \max_i p_i &= \frac{1 + (n-1)(1-\bar{\tau})}{n} = \psi, p = v_n(\psi) \end{aligned} \quad (12)$$

The closest point of the decision boundary to u_n is $v_n(1 + \bar{\tau}((1-n)/n))$ and hence on the line $\ell_{c_n, i}$ for each respective corner i . This implies, unlike uncertainty methods, proposed boundary does not interfere with the inference, (c) function. By definition (9) is robust to number of categories and the cases where one of the classes is already unfavorable. We omit derivations and refer the reader to [32].

Proposition 5. *Given $\tau, p = [p_1, p_2, \dots, p_n] \in \Delta_n$ s.t. $2 = \arg \max_{i \neq 1} p_i$ and $p_2 \geq (1-p_1)\tau$ and evidence $\varepsilon = [\varepsilon_+, 1, \dots, 1]$ where $\varepsilon_+ \sim \text{lognorm}(\mu, c^2)$, we define the*

posterior at sequence s $p^s = p^0 \oplus \varepsilon \oplus \dots \oplus \varepsilon$. With $\bar{\tau} = 2 - 2\tau$ and $\tilde{\tau} = ((2\tau - 1)(n - 1) + 1)/n$ define;

$$\begin{aligned} S_{r(M1)} &= \{p^s | p_1^s \geq \tau\} \\ S'_{R(M1)} &= \{p | p_2 \geq \tau\} \\ S_{R(MP)} &= \{p^s | p_1^s - p_2^s \geq 1 - \bar{\tau}\} \\ S'_{R(MP)} &= \{p^s | p_2^s - p_j^s \geq 1 - \bar{\tau}, j = \arg \max_{k \notin \{1,2\}} p_k^s\} \implies \\ p(p^s \in S_{R(MP)}(p_1, p_2, \bar{\tau}, s)) &\geq p(p^s \in S_{R(M1)}(p_1, \tau, s)) \text{ \& } \\ p(p^s \in S'_{R(M1)}(p_1, p_2, \tilde{\tau}, s)) &\geq p(p^s \in S'_{R(MP)}(p_1, p_2, \tilde{\tau}, s)) \\ &\geq p(p^s \in S'_{R(M1)}(p_1, p_2, \tau, s)) \forall s \in \{1, 2, \dots\} \end{aligned}$$

Proposition 5 demonstrates that compared to a stopping criterion based on posterior distribution thresholding (M1), the proposed method (MP) always has higher probability of entering the stopping region with correct decision, while the probability of entering an incorrect region resulting in incorrect decision is constrained within certain probability values. Therefore, a system can be designed to achieve a desired true positive probability while limiting false alarm probability by using the proposed stopping criterion. Supporting numerical examples are in the appendix Sec.7.3.1.

5 EXPERIMENTS AND RESULTS

In this section we run experiments to support our findings. Throughout this section we denote the *proposed method* introduced in Sec. 4 by (MP) and compare it with the methods presented in Table 1. Specifically, we designate a confidence level τ for (M1). We select respective c_s for (M2:4) such that (M1:4) intersect at $v_n(\tau)$ (See Observation 1). We select $\bar{\tau}$ for (MP) following Observation 4. For (M5) $c_{\delta_{KL}}$ is selected as 10^{-2} (see Table 1 for competing method definitions). We demonstrate results for the two following cases: (i) Synthetic experiments (ii) A letter decision for electroencephalography (EEG)-based brain computer interface (BCI) typing system.

5.1 Synthetic Experiments

In this section we aim to demonstrate the following items:

- 1) In the presence of disfavored classes, where the recursive classification is happening in a lower dimension than the cardinality of the class space (e.g. at least one of the class probabilities ≈ 0), uncertainty based methods (especially Renyi Entropy with orders $\alpha \leq 1$) suffer from immediate or rushed stopping (Sec. 3.1) whereas (MP) is robust to this effect and present similar performance as confidence based methods. (Details in Table 2, Table 3)
- 2) In case of a central posterior trajectory, confidence and uncertainty based methods perform similarly whereas (MP) stops earlier with marginal accuracy loss (Sec. 3.2). It is possible to early stop by lowering the confidence level. However, doing so is penalized in accuracy for cases in 1. (MP) provides an early stopping as if confidence level is relaxed for 2 and provides high accuracy performance for 1. (Details in Table 4)

In our experiments we present scenarios with pre-defined prior information; for Tables 2,3 the priors are set such that some of the classes have ~ 0 probability mass therefore this prior is close to the simplex edge and the posterior is

Method		Number of Sequences						
		1	2	3	4	5	6	7
MP	p_{stop}	0.00	0.06	0.22	0.43	0.59	0.67	0.77
	$p_{true stop}$	0.00	0.67	0.86	0.96	0.97	0.98	0.99
M1	p_{stop}	0.00	0.06	0.22	0.43	0.59	0.67	0.77
	$p_{true stop}$	0.00	0.67	0.86	0.96	0.97	0.98	0.99
M2	p_{stop}	0.00	0.08	0.27	0.47	0.62	0.70	0.79
	$p_{true stop}$	0.00	0.67	0.85	0.95	0.95	0.97	0.98
M3	p_{stop}	0.00	0.34	0.56	0.70	0.80	0.85	0.90
	$p_{true stop}$	0.00	0.20	0.48	0.63	0.74	0.80	0.87
M4	p_{stop}	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	$p_{true stop}$	0.00	0.55	0.72	0.82	0.87	0.90	0.94
M5	p_{stop}	0.00	0.00	0.37	0.41	0.46	0.55	0.66
	$p_{true stop}$	0.00	0.00	0.58	0.76	0.87	0.91	0.95
M1	p_{stop}	0.00	0.16	0.37	0.56	0.71	0.77	0.84
	$p_{true stop}$	0.00	0.64	0.84	0.93	0.95	0.97	0.98

TABLE 2. Accuracy of the proposed methods on the LION dataset. The best results are highlighted in bold.

$p(\text{cat}) = 0.42$
 $p(\text{lion}) = 0.55$
s.t. $|\mathcal{A}| = 3$

TABLE 2: A toy example with 3 class recursive classification. The set contains "cat, lion, dog" whereas the prior information for the given image (left) is $p = [0.42, 0.55, 0.03]$ with $\tau = .8$. Observer that the prior is located on the edge (See Sec. 3.1). We proceed in the recursive classification with evidence $\varepsilon = [\varepsilon_+, \varepsilon_-, \varepsilon_-]$ where $\varepsilon_+ \sim \text{lognorm}(0.6, 0.5^2)$ and $\varepsilon_- \sim \text{lognorm}(0, 0.5^2)$. We perform 5000 MC simulations to report each number. On the table (right), we present probability of stopping with the given criteria and accuracy among times the system stopped. As expected entropy (M3) staggers by stopping 34% and with only 20% accuracy within second sequence. More importantly, as we reasoned before, Renyi entropies with order $\alpha < 1$ (M4), $\alpha = 0.2$ in this case, create ambiguous decision regions which leads to stopping all the time as presented. (imagesrc:https://www.amazon.com/Pet-Krewe-PK00101-Costume-Small/dp/B010E4TAKW)

expected to move in a lower dimensional space (case 1). For Table 4 the prior is set such that the target probability is set to .1 where the rest of the class options share the remaining mass equally. This puts the prior close to the center of the opposing edge of the respective corner, therefore the posterior will follow a central path (case 2). In our experiments evidence (likelihood at each sequence) employs the vector form: $\varepsilon = [\varepsilon_+, \varepsilon_-, \dots, \varepsilon_-]$ where 1st candidate is the true class WLOG. In each sequence, In each sequence, following the assumptions of Fig. 1, ε_+ and all ε_- s each are sampled from their respective lognormal distributions independently (specifications are listed under each Table 2, 3, 4). The likelihoods are further merged with the latest posterior with $\oplus$ (3). For our experiments, in addition to the methods presented in Table 1, we also use a confidence lower bound with a confidence level $\tilde{\tau}$ that is derived in Observation 5 (e.g., $\tilde{\tau}$ for which the minimum accuracy can be achieved as (MP)) and we represent this method with ($\bar{M1}$). To be precise, for (M1) and (MP) presented in Fig. 4-(b), ($\bar{M1}$) will be the confidence lines each intersect with (MP) at the respective peak points. By design, we expect to show that (MP) is as robust as (M1) and (MP) reaches the speed of ($\bar{M1}$) with marginal accuracy decrease. In these tables, we report two measures for each method given number of sequences such that p_{stop} represents the probability of stopping, $p_{true|stop}$ represents the accuracy among all terminations (e.g. $p_{true|stop} = 0.5, p_{stop} = 0.6$ means the system correctly selected 300 out of 600 terminated in 1000 Monte Carlo (MC) simulations). In our results, we highlight $p_{stop} \geq .5$ indicating a reasonable termination operational point of (MP). E.g. in Table 2 (MP) at 5th sequence terminated 59% of all simulated classification tasks with 97% accuracy. Observe that high probability of termination p_{stop} is desired only if the accuracy $p_{true|stop}$ within the terminated recursions is high enough. Table 2 and Table 3 collectively explain items (1)-(2) such that, in these tables 1/3 and 7/10 classes are disfavored already in the beginning of estimation. Priors with disfavored classes lead to a trajectory that follows a path closer to sides of the simplex and hence result in rushed or immediate stopping



$p(\text{cat}) = 0.13$
 $p(\text{lion}) = 0.52$
 $p(\text{dog}) = 0.30$
 $\text{s.t. } |A| = 10$

Method		Number of Sequences								
		1	2	3	4	5	6	7	8	9
MP	p_{stop}	0.00	0.08	0.26	0.42	0.61	0.70	0.78	0.82	0.88
	$p_{\text{true stop}}$	0.00	0.36	0.78	0.89	0.94	0.96	0.97	0.98	0.98
M1	p_{stop}	0.00	0.03	0.18	0.35	0.54	0.66	0.76	0.81	0.86
	$p_{\text{true stop}}$	0.00	0.29	0.83	0.90	0.94	0.97	0.97	0.98	0.98
M2	p_{stop}	0.00	0.05	0.23	0.41	0.61	0.72	0.81	0.85	0.88
	$p_{\text{true stop}}$	0.00	0.32	0.80	0.89	0.94	0.96	0.97	0.98	0.98
M3	p_{stop}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	$p_{\text{true stop}}$	0.00	0.43	0.63	0.74	0.84	0.87	0.91	0.93	0.94
M4	p_{stop}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	$p_{\text{true stop}}$	0.00	0.43	0.63	0.74	0.84	0.87	0.91	0.93	0.94
M5	p_{stop}	0.00	0.00	0.35	0.40	0.47	0.55	0.62	0.69	0.76
	$p_{\text{true stop}}$	0.00	0.00	0.46	0.65	0.82	0.88	0.92	0.95	0.96
M1	p_{stop}	0.00	0.48	0.67	0.79	0.88	0.92	0.94	0.95	0.97
	$p_{\text{true stop}}$	0.00	0.38	0.70	0.80	0.87	0.90	0.93	0.95	0.96

TABLE 3: Toy example employing the properties listed in Table2. Here we are in a 10 class recursive classification scenario. To highlight the difference we set $\tau = .75$ and $\varepsilon_+ \sim \text{lognorm}(0.8, 0.5^2)$ and $\varepsilon_- \sim \text{lognorm}(-0.3, 0.5^2)$. We perform 5000 MC simulations to report each number. Here we start the recursive classification where 3 candidates share the probability weight and the remaining. Share the remainder almost uniformly. Observe that as class space expands (M3) and (M4) behave similarly. Most importantly (MP) is robust and behave similar to (M1) where the lower bound (M1) results a deficit in accuracy. (imagesrc:https://people.com/pets/tomsenn-lion-mane-dog-costume-amazon/)

Method		Number of Sequences							
		10	11	12	13	14	15	16	17
MP	p_{stop}	0.09	0.18	0.31	0.45	0.60	0.70	0.77	0.86
	$p_{\text{true stop}}$	0.95	0.98	0.98	0.99	0.99	0.99	1.00	1.00
M1	p_{stop}	0.05	0.11	0.21	0.33	0.48	0.60	0.70	0.81
	$p_{\text{true stop}}$	1.00	0.99	0.99	1.00	1.00	1.00	1.00	1.00
M2	p_{stop}	0.05	0.11	0.21	0.33	0.48	0.61	0.70	0.81
	$p_{\text{true stop}}$	1.00	0.99	0.99	1.00	1.00	1.00	1.00	1.00
M3	p_{stop}	0.05	0.12	0.22	0.34	0.49	0.61	0.71	0.81
	$p_{\text{true stop}}$	1.00	0.99	0.99	1.00	1.00	1.00	1.00	1.00
M4	p_{stop}	0.05	0.13	0.23	0.35	0.49	0.62	0.72	0.82
	$p_{\text{true stop}}$	0.98	0.99	0.99	1.00	1.00	1.00	1.00	1.00
M5	p_{stop}	0.00	0.00	0.00	0.00	0.00	0.01	0.03	0.08
	$p_{\text{true stop}}$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	1.00
M1	p_{stop}	0.10	0.20	0.33	0.47	0.61	0.72	0.79	0.88
	$p_{\text{true stop}}$	0.94	0.98	0.98	0.99	0.99	0.99	0.99	1.00

TABLE 4: Toy example employing the properties listed in Table2. Prior information in MC simulations has a fixed prior probability for the true class that is set to 0.1 and the remainder probability is distributed using a normalization to 0.9 sum value after random assignment at each simulation. Therefore the prior is in a central position away from the corner of true class and will follow a central path (See Sec.3.2). We set $\tau = .85$ and $\varepsilon_+ \sim \text{lognorm}(0.8, 0.5^2)$ and $\varepsilon_- \sim \text{lognorm}(-0.3, 0.5^2)$. Observe that (MP) provides early stopping achieving the speed of (M1) by losing marginal accuracy compared to (M1). As the posterior is following a central path, uncertainty based methods and confidence methods behave similarly as explained in Sec.3.2.

yielding probability $p(\text{stop}) \approx 1$ within few sequences. However, for both (M3)(M4) accuracy $p(\text{true|stop})$ is very low and as we increase dimensionality $3 \rightarrow 10$ in Table 3 (M3) characteristics gets closer to (M4) (as also discussed in the manuscript Fig. 3-(a)). In such cases a method should collect evidence to increase its accuracy to a desired level and as promised, (MP) performs comparably to confidence based methods reaching $p(\text{true|stop}) \geq .94$ in exchange for 5 sequences. Observe lowering confidence threshold with (M1) results in a significant 27% accuracy loss (Table 3) for which (MP) still achieves high accuracy. Table 4 demonstrates (3) where the prior is designed such that the probability of the true class is set to 0.1 where the remaining probabilities are almost equally distributed. This places the prior probability on a centric position and hence the posterior follows a more central path. Differently in this experiment all methods require ≥ 10 sequences to terminate. Following Sec. 3.2, uncertainty based methods and confidence behave similar. In such scenarios (MP) achieves similar accuracy $p(\text{true|stop})$ 1 sequence earlier at 14th sequence with $p(\text{stop}) \geq .5$ compared to other methods (i.e., (MP) is faster with the same accuracy compared to (M1-5)). Observe that, (MP) stops earlier at a lower confidence threshold similar to (M1) but this is without the accuracy expense in previous scenario where (M1) sacrifices 24% accuracy (see Table 3 (M1) column-3).




Fig. 5: Summary of operation characteristics of the methods using the starting conditions for recursive classification presented in Table3 and Table4 respectively on top and bottom. For our experiments we use $\varepsilon_+ \sim \text{lognorm}(0.6, 0.5^2)$ and $\varepsilon_- \sim \text{lognorm}(0, 0.5^2)$. We run 5000 recursive classification simulations and report average accuracies and average sequences for different τ values. Each line in the figure are drawn with a collection of (accuracy, sequence) points where each are computed for $\tau = [0.65, 0.69, 0.72, 0.76, 0.79, 0.83, 0.86, 0.9]$ from left to right. We omit (M5) as it spends way more sequences. Observe that (M4) staggers, especially for the disfavored case (left). (MP) allows us to select an operation point that favors accuracy in the disfavored and gains speed by losing marginal accuracy where the posterior is following a central path (right).

To summarize the findings, we visualize time accuracy trade-off for each method for the cases presented in Tables 3 and 4. These results are presented in Fig. 5. Top figure in Fig. 5, we observe that (MP) allows us to select (M1)'s sub-curve with high accuracy unlike (M1L=M1, M2,M3). From bottom illustration, where the posterior trajectory follows a central path, we observe that (MP) provides a better speed-accuracy operation curve compared to all other methods.

5.2 Real Data Experiments

5.2.1 Experimental Details

In our experiments we use a BCI typing system called RSVP Keyboard presented in Orhan's work [33] and the implementation BCIPy [34]. The system is visualized for the stimulus screen and an actual healthy participant performing a task in Fig. 6. System is tasked to correctly classify user's intended symbol in mind (σ). The system utilizes EEG as evidence (ε) and a language model that serves as a prior ($p(\sigma|\mathcal{H}_0)$). A sequence in this application refers to a series of rapid letter presentations to the user and collecting respective evidence. At the beginning of each sequence the system selects a predetermined number of most likely letters (Φ_s) with respect to latest posterior distribution $p(\sigma|\mathcal{H}_{s-1})$. Due to noisy observations (ε_s) at each sequence, the system relies on recursive evidence collection. Collected EEG evidence and prior information is fused in a Bayesian fashion ($p(\sigma|\mathcal{H}_1) = p(\sigma|\mathcal{H}_0) \oplus p(\varepsilon_0|\sigma, \Phi_0)$ see Appendix 7.1). In the signal p300 evoked potential presence corresponds to a positive response and absence corresponds to a negative response. Once intended symbol appears on the computer screen ($\sigma \in \Phi_s$), subject's recorded brain signal evokes a distinguishable response [35]. The system terminates




Fig. 6: EEG driven Rapid Serial Visual Presentation (RSVP) keyboard typing interface. (a) The stimuli is flashed in the middle of the screen while the user is informed with the text above. (b) The user is conducting copying the phrase task (multiple copy letter tasks). The user is informed about the required phrase. EEG is collected on top of the scalp non-invasively.

evidence collection once a stopping condition is met and in this section we show that (MP) provides early stopping with marginal accuracy penalty.

Data Collection: Ten healthy participants (six females), 20-35 years old were recruited under IRB-130107 protocol approved by Northeastern University. A DSI-24 Wearable Sensing EEG Headset was used for data acquisition, at a sampling rate of 300 Hz with active dry electrodes. All participants performed the calibration session containing 100 sequences; each sequence includes 5 trials; and one trial in each sequence is the target symbol which is displayed on the screen prior to each sequence (RSVP paradigm). A sequence contains randomly ordered ten symbols with a pre-defined target symbol. EEG is acquired from 16 channels using the International 10–20 configuration (Fp1, Fp2, F3, F4, Fz, Fc1, Fc2, Cz, P1, P2, C1, C2, Cp3, Cp4, P5, P6). Recorded EEG are used to learn class conditional EEG evidence distributions.

Pre-Processing: In EEG-BCIs the primary interest of filtering is to extract the signal of interest components [36], [37]. We first filter the EEG signal to remove drifts and artifact-related high frequency components with a band pass filter for [1,50]Hz. After filtering, EEG is windowed to extract the respective evidence at each channel for stimuli presentations. Time-windowed data from different EEG channels is usually concatenated to obtain the EEG feature vector that has a high dimension because of using a multi-channel measurement. Therefore, dimensionality reduction using ICA or PCA is also needed [36]. The system relies on reducing EEG time series into one dimensional feature vector. Filtered multi-channel EEG data time windows are passed through channel-wise principal component analysis where the outputs are concatenated to an intermediate feature vector. We assume in each class, feature vectors are drawn from a multivariate Gaussian distribution and hence Regularized discriminant analysis (regularized quadratic discriminant analysis [38]) is a plausible choice that results in one dimensional representation of the signal. Each positive and negative sample in the calibration EEG data is reduced to a single dimensional feature and positive and negative feature distributions are learned accordingly.

Experimental Task: For our experiments we use these distributions in Copy-Letter task such that each user’s data is used to type a target letter within a pre-determined phrase for multiple phrases [39]. More specifically, for each letter typing scenario the user is tasked to respond to the system, and the




Fig. 7: Number of recursion spent and accuracy plots for recursive classification in BCI typing system. Each scenario is generated using human-in-the-loop calibration data trained generative models. In each figure results are presented in ascending order of performance measures (area under receiver operation characteristics curve (AUC)). Number of queries in each recursion from top to bottom, 5, 10, 15 respectively. Legend covers methods from left to right and dots on the figures represent respective accuracy values. The users tried to type “A” without any language model (uniform prior information). Top to bottom legend order is from left to right for each block.

decision is made when the cumulative evidence matches the correct letter after multiple recursions. Therefore, in such a setting, the decision chance level is 0.03%. To make a decision the system recursively queries the user with multiple letter flashes. We designate the number of queries to be presented at each sequence to be $N \in \{15, 10, 5\}$. We present the results for two conditions. First one is a typing scenario with no language model (uniform prior information). And in the second scenario there exists a language model for the requested typing and the candidate letter is in top 16. The choice of two scenarios illustrated in the numerical results is to represent the following; (i) Class priors are uniform (in the 28-vertex simplex, RBC starts from the center of the simplex labeled as u in Fig. 2-(d)). (ii) The prior probability for the correct class label (desired letter of the user) is selected to be significantly lower than in the first case (in the 28-vertex simplex, RBC starts further away from the target vertex). These cases represent typical situations with uninformative prior and adversarial prior, additionally these are the challenging cases for RBC.

5.2.2 Results

In our experiments we categorize the users based on their calibration performances. The measure of performance is the area under receiver operating characteristics curve (AUC) of classification based on the features extracted during calibration. We specifically selected user data with AUC performances $\{0.67, 0.72, 0.76, 0.81, 0.84, 0.87\}$ to have a spectrum of ranging performances.

We visualize our findings in Fig. 7 and 8 respectively. It is observed that (MP) switches the operation point to a location




Fig. 8: The same implementation as described in Fig.7. The difference here is the subjects were supported by a language model trying to type "IT_O" given "IT_" and hence even though "O" is not the top letter it is one of the likely letters.

such that faster results are obtained with the cost of small amount of decreases in accuracy. To show the significance of the (MP), we present here a scenario that includes correctly typing 100 letters on a computer screen. We refer the reader to appendix Sec.7.3.2 for the complete results that allowed us to generate visualizations.

Comparing with the conventional method (M1), when uniform prior is used, (MP) outperforms (M1) in terms of speed, i.e., (MP) and (M1) complete the same task with 1735 and 1945 sequences respectively. Accordingly (MP) saves 210 sequences. This corresponds to saving $3(m)30(s)/32(m)$ lifetime during typing. Additionally, if a language model prior is used for the same task, (MP) still outperforms (M1), 1580 sequences vs 1728 sequences saving $2(m)24(s)/28(m)48(s)$ lifetime during typing. These reported amount of time are computed under the condition that at the end of the task 100 letters are completely correctly typed including corrections of the wrongly typed letters (i.e., email is completely correct). Saving time is very crucial for practical BCI typing as these systems are designed for individuals with limited speech and physical abilities. Therefore, fatigue and discomfort caused by the BCI system are important factors that significantly affect the BCI typing performance, and limiting the time to complete the tasks accurately in practice.

6 CONCLUSION


In this paper, our focus was on the analysis and design of stopping criterion for recursive Bayesian classification. Accordingly, we proposed a stopping condition design to overcome the sensitivity to number of classes of the existing uncertainty methods, while still providing an early stopping opportunity with marginal accuracy cost. Proposed method offers true positive probability rates above the true positive rates of confidence thresholding while limiting the false

alarm rate which we have shown analytically. We experimentally validated that proposed approach sustains an accuracy value where uncertainty methods are suffering while still allow for early stopping where confidence based methods struggle. It is possible to stop earlier with a lower confidence threshold to stop as early as proposed method where posterior follows a central trajectory. Synthetic experiments show that, lowering the confidence threshold is penalized in accuracy for disfavored cases where the proposed approach still sustained a high accuracy level. This emphasizes the importance of utilizing RBC geometry not to experience unwanted accuracy loss in certain scenarios. Therefore, it is possible to say the proposed method is combining the strengths of conventional methods. We also validated the practical use case on a brain computer interfaced typing system. This paper focuses on proposing a perspective and the proposed piece wise linear approach in Sec. 4 is a simplistic yet powerful design. We aim to explore the limitations of the curve order in stopping criterion design and propose new objectives that result in better performance. Also in Sec. 5 we experiment with evidence models with varying performances. A future work can center around providing an analytical relationship between evidence model performance (e.g. variance, AUC etc.) and stopping region performance. This work can further be extended by considering true positive maximization and false negative minimization as a multi-objective optimization.

REFERENCES

- [1] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*. John Wiley & Sons, 2012.
- [2] H. L. Van Trees and K. L. Bell, *Detection estimation and modulation theory*. Wiley, 2013.
- [3] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3.
- [4] A. M. Alaa and M. Van Der Schaar, "Balancing suspense and surprise: Timely decision making with endogenous information acquisition," in *NEURIPS*, 2016, pp. 2910–2918.
- [5] V. Vapnik, "Principles of risk minimization for learning theory," in *NEURIPS*, 1992, pp. 831–838.
- [6] G. Lugosi and K. Zeger, "Nonparametric estimation via empirical risk minimization," *IEEE Transactions on information theory*, vol. 41, no. 3, pp. 677–687, 1995.
- [7] M. Higger, F. Quivira, M. Akcakaya, M. Moghadamfalahi, H. Nezamfar, M. Cetin, and D. Erdogmus, "Recursive bayesian coding for bcis," *IEEE TNSRE*, vol. 25, no. 6, pp. 704–714, 2017.
- [8] S. Tong and D. Koller, "Support vector machine active learning with applications to text classification," *Journal of machine learning research*, vol. 2, no. Nov, pp. 45–66, 2001.
- [9] A. Wilson, A. Fern, and P. Tadepalli, "A bayesian approach for policy learning from trajectory preference queries," in *NEURIPS*, 2012, pp. 1133–1141.
- [10] T. Tsiligkaridis, B. M. Sadler, and A. O. Hero, "Collaborative 20 questions for target localization," *IEEE Transactions on Information Theory*, vol. 60, no. 4, pp. 2233–2252, 2014.
- [11] S. K. Jha, E. M. Clarke, C. J. Langmead, A. Legay, A. Platzter, and P. Zuliani, "A bayesian approach to model checking biological systems," in *International Conference on Computational Methods in Systems Biology*. Springer, 2009, pp. 218–234.
- [12] I. Sason and S. Verdú, "Arimoto–rényi conditional entropy and bayesian m -ary hypothesis testing," *IEEE Transactions on Information Theory*, vol. 64, no. 1, pp. 4–25, 2017.

- [13] S. Geisser, *Predictive inference*. Routledge, 2017.
- [14] D. Golovin, A. Krause, and D. Ray, "Near-optimal bayesian active learning with noisy observations," in *NEURIPS*, 2010, pp. 766–774.
- [15] A. N. Shiryaev, *Optimal stopping rules*. Springer Science & Business Media, 2007, vol. 8.
- [16] A. Irci, A. Saranlı, and B. Baykal, "Study on q-ram and feasible directions based methods for resource management in phased array radar systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 1848–1864, 2010.
- [17] A. Charlish, F. Hoffmann, R. Klemm, U. Nickel, and C. Gierull, "Cognitive radar management," *Novel radar techniques and applications: waveform diversity and cognitive radar and target tracking and data fusion*, vol. 2, pp. 157–193, 2017.
- [18] Y. Xiang, M. Akcakaya, S. Sen, D. Erdogmus, and A. Nehorai, "Target tracking via recursive bayesian state estimation in cognitive radar networks," *Signal Processing*, vol. 155, pp. 157–169, 2019.
- [19] Y. M. Marghi, A. Koçanaoğulları, M. Akçakaya, and D. Erdoğan, "A history-based stopping criterion in recursive bayesian state estimation," in *2019 ICASSP*. IEEE, 2019, pp. 3362–3366.
- [20] K. Bulatov, N. Razumnyi, and V. V. Arlazarov, "On optimal stopping strategies for text recognition in a video stream as an application of a monotone sequential decision model," *IJDAR*, vol. 22, no. 3, pp. 303–314, 2019.
- [21] T. L. Lai, "On optimal stopping problems in sequential hypothesis testing," *Statistica Sinica*, vol. 7, no. 1, pp. 33–51, 1997.
- [22] —, "Sequential analysis: some classical problems and new challenges," *Statistica Sinica*, pp. 303–351, 2001.
- [23] C. E. Shannon, "A mathematical theory of communication," *Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [24] Y. Li and R. E. Turner, "Rényi divergence variational inference," in *NEURIPS*, 2016, pp. 1073–1081.
- [25] D. Weinshall, H. Hermansky, A. Zweig, J. Luo, H. Jimison, F. Ohl, and M. Pavel, "Beyond novelty detection: Incongruent events, when general and specific classifiers disagree," in *NEURIPS*, 2009, pp. 1745–1752.
- [26] D. S. Pavlichin and T. Weissman, "Chained kullback-leibler divergences," in *2016 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2016, pp. 580–584.
- [27] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with bregman divergences," *Journal of machine learning research*, vol. 6, no. Oct, pp. 1705–1749, 2005.
- [28] V. Pawlowsky-Glahn and J. J. Egozcue, "Geometric approach to statistical analysis on the simplex," *Stochastic Environmental Research and Risk Assessment*, vol. 15, no. 5, pp. 384–398, 2001.
- [29] S.-i. Amari and H. Nagaoka, *Methods of information geometry*. American Mathematical Soc., 2007, vol. 191.
- [30] J. Aitchison, "The statistical analysis of compositional data," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 44, no. 2, pp. 139–160, 1982.
- [31] S.-W. Ho and S. Verdú, "On the interplay between conditional entropy and error probability," *IEEE Transactions on Information Theory*, vol. 56, no. 12, pp. 5930–5942, 2010.
- [32] J. Kittler and C. Zor, "Delta divergence: A novel decision cognizant measure of classifier incongruence," *IEEE transactions on cybernetics*, vol. 49, no. 6, pp. 2331–2343, 2018.
- [33] U. Orhan, K. E. Hild, D. Erdogmus, B. Roark, B. Oken, and M. Fried-Oken, "Rsvp keyboard: An eeg based typing interface," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 645–648.
- [34] T. Memmott, A. Koçanaoğulları, M. Lawhead, D. Klee, S. Dudy, M. Fried-Oken, and B. Oken, "Bcipy: Brain-computer interface software in python," *arXiv preprint arXiv:2002.06642*, 2020.
- [35] G. F. Woodman, "A brief introduction to the use of event-related potentials in studies of perception and attention," *Attention, Perception, & Psychophysics*, vol. 72, no. 8, pp. 2031–2046, 2010.
- [36] M. Akcakaya, B. Peters, M. Moghadamfalahi, A. Mooney, U. Orhan, B. Oken, D. Erdogmus, and M. Fried-Oken, "Noninvasive brain computer interfaces for augmentative and alternative communication," *IEEE Reviews in Biomedical Engineering*, vol. 7, no. 1, pp. 31–49, 2014.
- [37] M. Moghadamfalahi, U. Orhan, M. Akcakaya, H. Nezamfar, M. Fried-Oken, and D. Erdogmus, "Language-model assisted brain computer interface for typing: A comparison of matrix and rapid serial visual presentation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 5, pp. 910–920, 2015.
- [38] J. H. Friedman, "Regularized discriminant analysis," *Journal of the American statistical association*, vol. 84, no. 405, pp. 165–175, 1989.
- [39] U. Orhan, H. Nezamfar, M. Akcakaya, D. Erdogmus, M. Higger, M. Moghadamfalahi, A. Fowler, B. Roark, B. Oken, and M. Fried-Oken, "Probabilistic simulation framework for eeg-based bci design," *Brain-Computer Interfaces*, vol. 3, no. 4, pp. 171–185, 2016.
- [40] C. Barceló-Vidal, J. A. Martín-Fernández, and V. Pawlowsky-Glahn, "Mathematical foundations of compositional data analysis," in *Proceedings of IAMG*, vol. 1, 2001, pp. 1–20.
- [41] A. Rényi *et al.*, "On measures of entropy and information," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California, 1961.



Aziz Koçanaoğulları received the B.S. degree in electronics engineering and mathematics from Istanbul Technical University, Istanbul Turkey in 2014 and 2015 respectively. He received M.Sc. degree in telecommunications engineering in Istanbul Technical University in 2016. He received Ph.D. degree in electrical engineering in Northeastern University Boston, MA, USA in 2020. Since then he is a member of Computational Radiology Laboratory (CRL) in Boston Children's Hospital - Harvard Medical School Boston, MA, USA where he is currently a post doctoral research fellow. His main areas of research interest are active recursive inference in sequential decision making processes and active model learning.



Murat Akcakaya received the Ph.D. degree in electrical engineering from Washington University in St. Louis, MO, USA, in December 2010. He is an Associate Professor in the Electrical and Computer Engineering Department of the University of Pittsburgh. His research interests are in the areas of statistical signal processing and machine learning.



Deniz Erdoğan received the B.S. degree in electrical engineering and mathematics, and the M.S. degree in electrical engineering from the Middle East Technical University, Ankara, Turkey, in 1997 and 1999, respectively, and the Ph.D. degree in electrical and computer engineering, in 2002, from the University of Florida, Gainesville, FL, USA, where he was a postdoc until 2004. He is currently a Research Professor at Northeastern University. His research focuses on statistical signal processing and machine learning with applications to biomedical signal/image processing and cyberhuman systems.