

Article

Simple Stopping Criteria for Information Theoretic Feature Selection

Shujian Yu  and José C. Príncipe * 

Computational NeuroEngineering Laboratory, University of Florida, Gainesville, FL 32611, USA;
yusjlc9011@ufl.edu

* Correspondence: principe@cnel.ufl.edu; Tel.: +1-352-392-2682

Received: 4 December 2018; Accepted: 21 January 2019; Published: 21 January 2019



Abstract: Feature selection aims to select the smallest feature subset that yields the minimum generalization error. In the rich literature in feature selection, information theory-based approaches seek a subset of features such that the mutual information between the selected features and the class labels is maximized. Despite the simplicity of this objective, there still remain several open problems in optimization. These include, for example, the automatic determination of the optimal subset size (i.e., the number of features) or a stopping criterion if the greedy searching strategy is adopted. In this paper, we suggest two stopping criteria by just monitoring the conditional mutual information (CMI) among groups of variables. Using the recently developed multivariate matrix-based Rényi's α -entropy functional, which can be directly estimated from data samples, we showed that the CMI among groups of variables can be easily computed without any decomposition or approximation, hence making our criteria easy to implement and seamlessly integrated into any existing information theoretic feature selection methods with a greedy search strategy.

Keywords: feature selection; stopping criterion; conditional mutual information; multivariate matrix-based Rényi's α -entropy functional

1. Introduction

Feature selection aims to find the smallest feature subset that yields the minimum generalization error [1]. As a fundamental problem in machine learning and statistics communities, feature selection is relevant to understand better the classification problem at hand in many applications, such as medical decision making, economics, and engineering [2]. Consequently, it also has a large impact on interpretability or explainability, which is totally missing in most current deep learning techniques [3]. Ever since the pioneering work of Battiti [4], information theoretic feature selection has been extensively investigated in signal processing and machine learning communities (e.g., [5–7]). Given a set of F features $S = \{x_1, x_2, \dots, x_F\}$ (each x_i denotes an attribute) and their corresponding class labels y , these methods seek a subset of informative attributes $S^* \subset S$, such that the mutual information (MI) between S^* and y (i.e., $\mathbf{I}(S^*; y)$) is maximized [8].

Despite the simplicity of this objective, several open problems still remain in information theoretic feature selection. These include, for example, the reliable estimation of $\mathbf{I}(S'; y)$ in high-dimensional space, in which S' denotes an arbitrary subset of S [8,9]. In fact, S' may contain both continuous and discrete variables, whereas y is a discrete variable. There is no universal agreement on the definition of MI between a discrete variable and a group of mixed variables, let alone its estimation [10]. Therefore, almost all existing information theoretic feature selection methods estimate $\mathbf{I}(S'; y)$ by first discretizing the feature space and then approximating $\mathbf{I}(S'; y)$ with low-order MI quantities, such as relevancy $\mathbf{I}(x_i; y)$, joint relevancy $\mathbf{I}(\{x_i, x_j\}; y)$, conditional relevancy $\mathbf{I}(x_i; y|x_j)$, redundancy $\mathbf{I}(x_i; x_j)$, conditional redundancy $\mathbf{I}(x_i; x_j|y)$, and synergy [11]. These low-order MI quantities only capture the low-order

feature dependency and, hence, severely limit the performance of existing information theoretic feature selection methods [12]. Interested readers can check Reference [8] for a systemic review of 17 popular low-order information theoretic criteria in the last two decades. Apart from MI estimation, another challenging problem, which also impacts performance greatly, is the automatic determination of the optimal size of S^* . This is because most of the information theoretic feature selection methods do not have a stopping criterion [1]. Hence, the user must find criteria to estimate the best number of features.

Regarding the first problem, our recent work [13] suggested that $\mathbf{I}(S'; y)$ can be efficiently estimated using the normalized eigenspectrum of a Hermitian matrix of the projected data in the reproducing kernel Hilbert space (RKHS). In this paper, we expand on the topic of Reference [13] and illustrate that the novel multivariate matrix-based Rényi's α -entropy functional also enables simple strategies to guide the early stopping in the greedy search procedure of information theoretic feature selection methods.

Before presenting our methods, we first briefly review related work and also point out a common issue in most previous methods. Perhaps the most acknowledged stopping criterion for information theoretic feature selection is that the value of $\mathbf{I}(S'; y)$ stops increasing or reaches its maximum [14,15]. Given a new feature x' , this rule suggests that we should stop selection if $\mathbf{I}(\{S', x'\}; y) \leq \mathbf{I}(S'; y)$. An alternative approach is using the concept of the Markov blanket (MB) [16,17]. By definition, the MB M of a target variable y is the smallest subset of S such that y is conditional independent of the rest of the variables $S - M$, i.e., $y \perp (S - M) | M$ [1]. From the perspective of information theory, this rule indicates that we should stop selection if the conditional mutual information (CMI) $\mathbf{I}(\{S - M\}; y | M)$ is zero.

Despite their simplicity, both rules are overoptimistic and cannot be applied in practice. In fact, by the chain rule of mutual information [18], we have:

$$\mathbf{I}(\{S', x'\}; y) = \mathbf{I}(S'; y) + \mathbf{I}(x'; y | S'), \quad (1)$$

and:

$$\mathbf{I}(\{S - M\}; y | M) = \mathbf{I}(S; y) - \mathbf{I}(M; y). \quad (2)$$

According to Equations (1) and (2), the feature selection is stopped if and only if the CMI item $\mathbf{I}(x'; y | S')$ or $\mathbf{I}(\{S - M\}; y | M)$ is zero. Unfortunately, since CMI is always non-negative [18] and rarely reduces to zero in practice due to statistical variation and chance agreement between variables [19], we always have $\mathbf{I}(\{S', x'\}; y) > \mathbf{I}(S'; y)$ and $\mathbf{I}(\{S - M\}; y | M) > 0$. That is to say, the maximum value of $\mathbf{I}(S'; y)$ is exactly $\mathbf{I}(S; y)$ and a perfect MB of y is perhaps the feature set S itself.

Admittedly, one can say that we can stop the selection if the increment of $\mathbf{I}(S'; y)$ or the decrement of $\mathbf{I}(\{S - S'\}; y | S')$ approaches zero with a tiny residual. Unfortunately, since we still do not have a reliable estimator of MI and CMI in high-dimensional space (before [13]), it is hard for us to measure or determine how small the residual terms are.

To the best of our knowledge, there are only two methods in the literature that can stop the greedy search. François et al. [14] suggested monitoring the value of $\mathbf{I}(S'; y)$ using a permutation test [20]. Specifically, suppose the new feature selected in the current iteration is x^{cand} , the authors create a random permutation of x^{cand} (without permuting the corresponding y), denoted \tilde{x}^{cand} . If $\mathbf{I}(\{S', x^{\text{cand}}\}; y)$ is not significantly larger than $\mathbf{I}(\{S', \tilde{x}^{\text{cand}}\}; y)$, x^{cand} can be discarded and the feature selection is stopped. Moreover, François et al. also suggested using the k -nearest neighbors (KNN) estimator [21] to estimate $\mathbf{I}(S'; y)$. However, their results indicated that this estimator may result in negative CMI quantities no matter the selection of k . Vinh et al. [19], on the other hand, proposed monitoring the increment of $\mathbf{I}(S'; y)$ after adding x^{cand} (i.e., $\mathbf{I}(x^{\text{cand}}; y | S')$) using χ^2 distribution. If $\mathbf{I}(x^{\text{cand}}; y | S')$ is smaller than a threshold obtained from the χ^2 distribution at a certain significance level, the feature selection is stopped. However, one should note that the χ^2 distribution assumption holds if and only if $\mathbf{I}(\tilde{x}^{\text{cand}}; y | S') = 0$ [19,22]. Unfortunately, as we emphasized earlier, the CMI

quantity rarely reduces to zero in practice. This is also the reason why Vinh's method is likely to severely underestimate the required number of features, as is illustrated in the experiments.

Different from previous work, we suggest using the novel multivariate matrix-based Rényi's α -entropy functional [23] to estimate MI and CMI in high-dimensional space. We also present two simple stopping criteria based on these new estimators. To summarize, our main contributions are twofold:

- Efficient stopping criteria for feature selection are still missing in the literature. We present two rules based on the new estimators of information theoretic descriptors of Rényi's α -entropy in RKHS. One should note that the properties, utilities, and possible applications of these new estimators are rather new and mostly unknown to practitioners;
- Our criteria are extremely simple and flexible. First, the proposed tests and stopping criteria can be incorporated into any sequential feature selection procedure, such as mutual information-based feature selection (MIFS) [4] and maximum-relevance minimum-redundancy (MRMR) [24]. Second, benefiting from the higher accuracy of the estimation, we demonstrated that one can simply use a threshold in a heuristic manner, which is very valuable for practitioners across different domains.

2. Simple Stopping Criteria for Information Theoretic Feature Selection

In this section, we start with a brief introduction to the recently proposed matrix-based Rényi's α -entropy functional and its multivariate extension [13]. The novel definition yields two simple stopping criteria as presented below.

2.1. Matrix-Based Rényi's α -Entropy Functional and Its Multivariate Extension

In information theory, a natural extension of the well-known Shannon's entropy is Rényi's α -order entropy [25]. For a random variable X with probability density function (PDF) $f(x)$ in a finite set \mathcal{X} , the α -entropy $\mathbf{H}_\alpha(X)$ is defined as:

$$\mathbf{H}_\alpha(f) = \frac{1}{1-\alpha} \log \int_{\mathcal{X}} f^\alpha(x) dx. \quad (3)$$

Based on this entropy definition, Rényi then proposed a divergence measure (α -relative entropy) between random variables with PDFs f and g :

$$\mathbf{D}_\alpha(f||g) = \frac{1}{\alpha-1} \log \int_{\mathcal{X}} f(x) \left(\frac{f(x)}{g(x)} \right)^{\alpha-1} dx. \quad (4)$$

Rényi's entropy and divergence have a long track record of usefulness in information theory and its applications [26]. Unfortunately, the accurate PDF estimation impedes its more widespread adoption in data driven science. To solve this problem, References [13,23] suggest similar quantities that resemble quantum Rényi's entropy [27] in terms of the normalized eigenspectrum of the Hermitian matrix of the projected data in RKHS, thus estimating the entropy and joint entropy among two or multiple variables directly from data without PDF estimation. For brevity, we directly give the definition.

Definition 1. Let $\kappa : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ be a real valued positive definite kernel that is also infinitely divisible [28]. Given $X = \{x^1, x^2, \dots, x^n\}$ and the Gram matrix K obtained from κ on all pairs of exemplars, that is, $(K)_{ij} = \kappa(x^i, x^j)$, a matrix-based analogue to Rényi's α -entropy for X can be defined, using a normalized positive definite (NPD) matrix A of size $n \times n$ and trace 1, by:

$$\mathbf{S}_\alpha(A) = \frac{1}{1-\alpha} \log_2 (\text{tr}(A^\alpha)) = \frac{1}{1-\alpha} \log_2 \left[\sum_{i=1}^n \lambda_i(A)^\alpha \right], \quad (5)$$

where $A_{ij} = \frac{1}{n} \frac{K_{ij}}{\sqrt{K_{ii}K_{jj}}}$ and $\lambda_i(A)$ denotes the i -th eigenvalue of A .

The matrix functional in Equation (5) bears a lot of resemblance to well-known operational quantities from quantum information theory [27], where the density matrix (operator) ρ can be employed to compute expectation over an observable represented by the operator A as $\langle A \rangle = \text{tr}(\rho A)$. For instance, the quantum extensions of Rényi’s entropy [29,30] is given by:

$$\mathbf{S}_\alpha(\rho) = \frac{1}{1-\alpha} \log [\text{tr}(\rho^\alpha)] . \tag{6}$$

Although some properties of Equation (6) also apply to Equation (5), the matrix-based estimation is very different, since it deals with the Gram matrix obtained from pairwise evaluations of a positive definite kernel from data samples. Consequently, this novel definition not only involves the functional, but also the kernels employed to construct positive definite matrix.

Definition 2. Given a collection of n samples $\{s_i = (x_1^i, x_2^i, \dots, x_k^i)\}_{i=1}^n$, where the superscript i denotes the sample index, each sample contains k ($k \geq 2$) measurements $x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2, \dots, x_k \in \mathcal{X}_k$ obtained from the same realization, and the positive definite kernels $\kappa_1 : \mathcal{X}_1 \times \mathcal{X}_1 \mapsto \mathbb{R}, \kappa_2 : \mathcal{X}_2 \times \mathcal{X}_2 \mapsto \mathbb{R}, \dots, \kappa_k : \mathcal{X}_k \times \mathcal{X}_k \mapsto \mathbb{R}$, let us denote $X_l = \{x_1^l, x_2^l, \dots, x_l^l\} (1 \leq l \leq k)$ containing the l -th measurement of all samples, the Rényi’s α -order joint-entropy among k variables (denote here $\mathbf{J}_\alpha(X_1, X_2, \dots, X_k)$) can be defined using the following matrix-based analogue:

$$\mathbf{J}_\alpha(X_1, X_2, \dots, X_k) = \mathbf{S}_\alpha \left(\frac{A_1 \odot A_2 \odot \dots \odot A_k}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_k)} \right), \tag{7}$$

where $(A_1)_{ij} = \kappa_1(x_1^i, x_1^j), (A_2)_{ij} = \kappa_2(x_2^i, x_2^j), \dots, (A_k)_{ij} = \kappa_k(x_k^i, x_k^j)$, and \odot denotes the Hadamard product.

The following two corollaries (proved in Reference [13]) serve as a foundation for our Definition 2. Specifically, Corollary 1 indicates that the joint entropy of a set of variables is greater than or equal to the maximum of all of the individual entropies of the variables in the set, whereas Corollary 2 suggests that the joint entropy of a set of variable is less than or equal to the sum of the individual entropies of the variables in the set.

Corollary 1. Let $[C]$ be the index set $\{1, 2, \dots, C\}$. We partition $[C]$ into two complementary subsets \mathbf{s} and $\tilde{\mathbf{s}}$. For any $\mathbf{s} \subset [C]$, denote all indices in \mathbf{s} with $s_1, s_2, \dots, s_{|\mathbf{s}|}$, where $|\cdot|$ stands for cardinality. Similarly, denote all indices in $\tilde{\mathbf{s}}$ with $\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_{|\tilde{\mathbf{s}}|}$. Also let A_1, A_2, \dots, A_C be C $n \times n$ positive definite matrices with trace 1 and non-negative entries, and $(A_1)_{ii} = (A_2)_{ii} = \dots = (A_C)_{ii} = \frac{1}{n}$, for $i = 1, 2, \dots, n$. Then, the following two inequalities hold:

$$\mathbf{S}_\alpha \left(\frac{A_1 \odot A_2 \odot \dots \odot A_C}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_C)} \right) \leq \mathbf{S}_\alpha \left(\frac{A_{s_1} \odot A_{s_2} \odot \dots \odot A_{s_{|\mathbf{s}|}}}{\text{tr}(A_{s_1} \odot A_{s_2} \odot \dots \odot A_{s_{|\mathbf{s}|}})} \right) + \mathbf{S}_\alpha \left(\frac{A_{\tilde{s}_1} \odot A_{\tilde{s}_2} \odot \dots \odot A_{\tilde{s}_{|\tilde{\mathbf{s}}|}}}{\text{tr}(A_{\tilde{s}_1} \odot A_{\tilde{s}_2} \odot \dots \odot A_{\tilde{s}_{|\tilde{\mathbf{s}}|}})} \right), \tag{8}$$

$$\mathbf{S}_\alpha \left(\frac{A_1 \odot A_2 \odot \dots \odot A_C}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_C)} \right) \geq \max \left[\mathbf{S}_\alpha \left(\frac{A_{s_1} \odot A_{s_2} \odot \dots \odot A_{s_{|\mathbf{s}|}}}{\text{tr}(A_{s_1} \odot A_{s_2} \odot \dots \odot A_{s_{|\mathbf{s}|}})} \right), \mathbf{S}_\alpha \left(\frac{A_{\tilde{s}_1} \odot A_{\tilde{s}_2} \odot \dots \odot A_{\tilde{s}_{|\tilde{\mathbf{s}}|}}}{\text{tr}(A_{\tilde{s}_1} \odot A_{\tilde{s}_2} \odot \dots \odot A_{\tilde{s}_{|\tilde{\mathbf{s}}|}})} \right) \right]. \tag{9}$$

Corollary 2. Let A_1, A_2, \dots, A_C be C $n \times n$ positive definite matrices with trace 1 and non-negative entries, and $(A_1)_{ii} = (A_2)_{ii} = \dots = (A_C)_{ii} = \frac{1}{n}$, for $i = 1, 2, \dots, n$. Then, the following two inequalities hold:

$$\mathbf{S}_\alpha \left(\frac{A_1 \odot A_2 \odot \dots \odot A_C}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_C)} \right) \leq \mathbf{S}_\alpha(A_1) + \mathbf{S}_\alpha(A_2) + \dots + \mathbf{S}_\alpha(A_C), \tag{10}$$

$$\mathbf{S}_\alpha \left(\frac{A_1 \odot A_2 \odot \dots \odot A_C}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_C)} \right) \geq \max[\mathbf{S}_\alpha(A_1), \mathbf{S}_\alpha(A_2), \dots, \mathbf{S}_\alpha(A_C)]. \tag{11}$$

2.2. Stopping Criteria Based on Conditional Mutual Information

Denote $\{x_1, x_2, \dots, x_k\}$, the selected features in S' , $\{x'_1, x'_2, \dots, x'_m\}$, the remaining features in $S - S'$, given the entropy and joint entropy estimators shown in Equations (5)–(7), the MI between y and S' (i.e., $\mathbf{I}(S'; y)$) and the CMI between y and $S - S'$ conditioning on S' (i.e., $\mathbf{I}(\{S - S'\}; y|S')$), with analogue properties to Shannon’s definition of MI and CMI (by Shannon’s definition, $\mathbf{I}(S'; y) = \mathbf{H}(S') + \mathbf{H}(y) - \mathbf{H}(S', y)$ and $\mathbf{I}(\{S - S'\}; y|S') = \mathbf{H}(S - S', S') + \mathbf{H}(y, S') - \mathbf{H}(S - S', y, S') - \mathbf{H}(S')$, where \mathbf{H} denotes entropy or joint entropy), which can be estimated with Equations (12) and (13), respectively, where $A_1, A_2, \dots, A_k, B, C_1, C_2, \dots, C_m$ denote the Gram matrices evaluated over $x_1, x_2, \dots, x_k, y, x'_1, x'_2, \dots, x'_m$, respectively, and \odot denotes the Hadamard product.

$$\mathbf{I}_\alpha(B; \{A_1, A_2, \dots, A_k\}) = \mathbf{S}_\alpha(B) + \mathbf{S}_\alpha\left(\frac{A_1 \odot A_2 \odot \dots \odot A_k}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_k)}\right) - \mathbf{S}_\alpha\left(\frac{A_1 \odot A_2 \odot \dots \odot A_k \odot B}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_k \odot B)}\right), \quad (12)$$

$$\begin{aligned} \mathbf{I}_\alpha(\{C_1, C_2, \dots, C_m\}; B|\{A_1, A_2, \dots, A_k\}) &= \mathbf{S}_\alpha\left(\frac{A_1 \odot A_2 \odot \dots \odot A_k \odot C_1 \odot C_2 \odot \dots \odot C_m}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_k \odot C_1 \odot C_2 \odot \dots \odot C_m)}\right) \\ &\quad + \mathbf{S}_\alpha\left(\frac{B \odot A_1 \odot A_2 \odot \dots \odot A_k}{\text{tr}(B \odot A_1 \odot A_2 \odot \dots \odot A_k)}\right) \\ &\quad - \mathbf{S}_\alpha\left(\frac{A_1 \odot A_2 \odot \dots \odot A_k \odot B \odot C_1 \odot C_2 \odot \dots \odot C_m}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_k \odot B \odot C_1 \odot C_2 \odot \dots \odot C_m)}\right) - \mathbf{S}_\alpha\left(\frac{A_1 \odot A_2 \odot \dots \odot A_k}{\text{tr}(A_1 \odot A_2 \odot \dots \odot A_k)}\right). \end{aligned} \quad (13)$$

As can be seen, the multivariate matrix-based Rényi’s α -entropy functional enables simple estimation of both MI and CMI in high-dimensional space, no matter the data characteristics (e.g., continuous or discrete) in each dimension. In contrast to previous definitions on Rényi’s α -mutual information [31] that stem from a probabilistic perspective using expectations over PDFs, Equations (12) and (13) directly estimate MI and CMI in RKHS. Benefiting from these elegant expressions, suppose the new feature selected in the current iteration is x^{cand} , we present two simple criteria to guide the early stopping of the greedy search. Specifically, we aim to test the “goodness-of-fit” of the MB condition, i.e., $S - S'$ is the MB of y given S' . Intuitively, if $\mathbf{I}(\{S - S' - x^{\text{cand}}\}; y|\{S', x^{\text{cand}}\})$ approaches zero, the MB condition is approximately satisfied.

Criterion I. If $\mathbf{I}(\{S - S' - x^{\text{cand}}\}; y|\{S', x^{\text{cand}}\}) \leq \varepsilon$, where ε refers to a tiny threshold, then we should stop the selection. We term this criterion CMI-heuristic, since ε is a heuristic value.

Criterion II. Motivated by Reference [14], in order to quantify how x^{cand} affects the MB condition, we created a random permutation of x^{cand} (without permuting the corresponding y), denoted \tilde{x}^{cand} . If $\mathbf{I}(\{S - S' - x^{\text{cand}}\}; y|\{S', x^{\text{cand}}\})$ is not significantly smaller than $\mathbf{I}(\{S - S' - \tilde{x}^{\text{cand}}\}; y|\{S', \tilde{x}^{\text{cand}}\})$, x^{cand} can be discarded and the feature selection is stopped. We term this criterion CMI-permutation (see Algorithm 1 for more details, where $\mathbb{1}$ is an indicator function which is 1 if its argument is true and 0 otherwise.).

Algorithm 1 CMI-permutation

Input: Feature set S ; Selected feature subset S' ; Class labels y ; Selected feature in the current iteration

x^{cand} ; Permutation number P ; Significance level θ .

Output: *decision* (Stop selection or Continue selection).

- 1: Estimate $\mathbf{I}(\{S - S' - x^{\text{cand}}\}; y | \{S', x^{\text{cand}}\})$ with Equation (13).
- 2: **for** $i = 1$ to P **do**
- 3: Randomly permute x^{cand} to obtain $\tilde{x}_i^{\text{cand}}$.
- 4: Estimate $\mathbf{I}(\{S - S' - \tilde{x}_i^{\text{cand}}\}; y | \{S', \tilde{x}_i^{\text{cand}}\})$ with Equation (13).
- 5: **end for**
- 6: **if** $\frac{\sum_{i=1}^P \mathbb{1}[\mathbf{I}(\{S - S' - x^{\text{cand}}\}; y | \{S', x^{\text{cand}}\}) \geq \mathbf{I}(\{S - S' - \tilde{x}_i^{\text{cand}}\}; y | \{S', \tilde{x}_i^{\text{cand}}\})]}{P} \leq \theta$ **then**
- 7: *decision* \leftarrow Continue selection.
- 8: **else**
- 9: *decision* \leftarrow Stop selection.
- 10: **end if**
- 11: **return** *decision*

3. Experiments and Discussions

We compared our two criteria with existing ones [14,19] on 10 well-known public datasets used in previous feature selection research [8,19], covering a wide variety of sample-feature ratios and a range of multiclass problems. The detailed properties of these datasets, including the number of features (F), the number of examples (N), and the number of classes (C), are available in Reference [8]. We refer the criterion in Reference [19] $\Delta\text{MI}-\chi^2$, since it monitors the increment of MI $\mathbf{I}(S'; y)$ (i.e., $\mathbf{I}(x^{\text{cand}}; y | S')$) with χ^2 distribution. We refer the criterion in Reference [14] MI-permutation, since it uses the permutation test to quantify the impact of x^{cand} on $\mathbf{I}(S'; y)$. Throughout this paper, we selected $\varepsilon = 10^{-4}$ in CMI-heuristic and used the multivariate matrix-based Rényi's α -entropy functional to estimate all MI quantities in MI-permutation. The baseline information feature selection method used in this paper is from Reference [13], which directly optimizes $\mathbf{I}(S'; y)$ in a greedy manner without any decomposition or approximation. An example for different stopping criteria on the dataset waveform is shown in Figure 1.

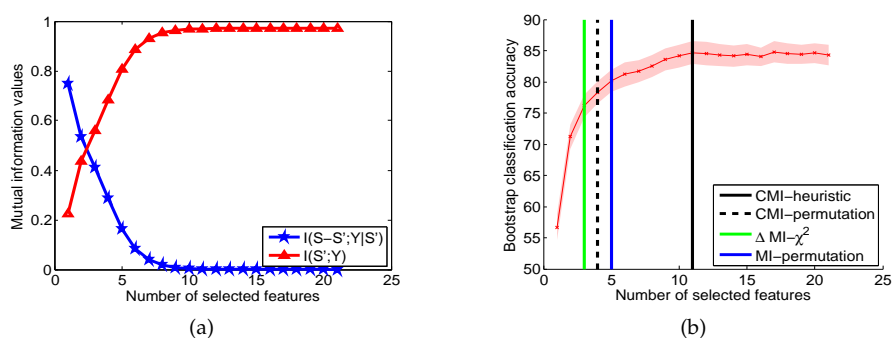


Figure 1. (a) shows the the values of mutual information (MI) $\mathbf{I}(S'; y)$ and conditional mutual information (CMI) $\mathbf{I}(\{S - S'\}; y | S')$ with respect to different number of selected features, i.e., the size of S' . $\mathbf{I}(S'; y)$ is monotonically increasing, whereas $\mathbf{I}(\{S - S'\}; y | S')$ is monotonically decreasing. (b) shows the terminated points produced by different stopping criteria, namely CMI-heuristic (black solid line), CMI-permutation (black dashed line), $\Delta\text{MI}-\chi^2$ (green solid line), and MI-permutation (blue solid line). The red curve with the shaded area indicates the average bootstrap classification accuracy with 95% confidence interval. In this example, the bootstrap classification accuracy reaches its statistical maximum value with 11 features and CMI-heuristic performs the best.

To provide a comprehensive evaluation, we tested the performance of the novel Rényi’s α -entropy estimator with $\alpha = 1.01$ (to approximate Shannon’s definition) and $\alpha = 2$ (i.e., the quadratic information quantities that have been widely applied in information theoretic learning [26]). Moreover, we also tested the performance of different statistical tests with significance level $\theta = 0.95$ and $\theta = 0.99$, but since the results are very similar, herein, we only report results with $\theta = 0.95$. The quantitative results, under different values of α , are summarized in Tables 1 and 2. For each criterion, we reported the number of selected features and the average classification accuracy across 100 bootstrap runs. In each run, N bootstrap samples are drawn for the training set, while the unselected samples serve as the test set. Same as in Reference [8], we used the linear support vector machine (SVM) [32] as the baseline classifier. To give a reference, we defined the “optimal” number of features (an unknown parameter) as the one that yields the maximum bootstrap accuracy or first achieves a bootstrap accuracy with no statistical difference to the maximum value (evaluated by paired t-test with significance level 0.05), and ranked all the criteria based on the difference between their estimated number of features and the optimal one.

Table 1. The number of selected features (#F) and the bootstrap classification accuracy (acc.) comparison for CMI-heuristic and CMI-permutation against different stopping criteria. We set $\alpha = 1.01$ for Rényi’s α -entropy functional. All criteria are ranked based on the difference between their selected number of features and the optimal values. The best two ranks are marked with green and blue, respectively. The average rank across all datasets is reported in the bottom line. The value behind the name of each dataset indicates the total number of features.

	CMI-Heuristic			CMI-Permutation			$\Delta MI-\chi^2$ [19]			MI-Permutation [14]			“Optimal”	
	#F	acc.	rank	#F	acc.	rank	#F	acc.	rank	#F	acc.	rank	#F	acc.
waveform (21)	11	84.7 ± 1.8	1	4	78.3 ± 1.8	3	3	76.1 ± 1.8	4	5	80.2 ± 1.8	2	11	84.7 ± 1.8
breast (30)	2	92.3 ± 1.7	1	2	92.3 ± 1.7	1	2	92.3 ± 1.7	1	2	92.3 ± 1.7	1	2	95.2 ± 1.2
heart (13)	13	81.7 ± 3.5	4	4	80.4 ± 3.3	1	2	76.9 ± 3.8	3	4	80.4 ± 3.3	1	6	82.6 ± 3.0
spect (22)	22	80.6 ± 3.6	4	11	82.1 ± 3.3	1	1	80.1 ± 3.3	3	7	81.1 ± 3.3	2	11	82.1 ± 3.3
ionosphere (34)	15	83.3 ± 2.8	1	7	81.8 ± 2.8	2	1	76.7 ± 3.2	4	7	81.8 ± 2.8	2	33	85.3 ± 3.0
parkinsons (22)	12	85.2 ± 3.7	1	4	85.0 ± 3.2	2	1	85.1 ± 3.5	4	4	85.0 ± 3.2	2	9	86.5 ± 3.4
semeion (256)	59	86.1 ± 1.3	1	20	77.7 ± 1.5	2	4	49.6 ± 1.7	4	20	77.7 ± 1.5	2	73	93.3 ± 1.3
Lung (325)	5	74.2 ± 7.7	3	10	73.9 ± 8.0	2	1	46.5 ± 7.5	4	13	79.1 ± 7.9	1	41	84.3 ± 6.5
Lymph (4026)	6	81.3 ± 5.8	1	248	88.7 ± 6.1	3	2	62.8 ± 6.5	2	249	88.9 ± 6.2	4	70	90.7 ± 5.4
Madelon (500)	3	69.5 ± 1.6	2	2	59.5 ± 1.6	3	4	76.7 ± 1.5	1	2	59.5 ± 1.6	3	4	76.7 ± 1.5
average rank			1.9			2.0			3.0			2.0		

Table 2. The number of selected features (#F) and the bootstrap classification accuracy (acc.) comparison for CMI-heuristic and CMI-permutation against different stopping criteria. We set $\alpha = 2$ for Rényi’s α -entropy functional. All criteria are ranked based on the difference between their selected number of features and the optimal values. The best two ranks are marked with green and blue, respectively. The average rank across all datasets is reported in the bottom line. The value behind the name of each dataset indicates the total number of features.

	CMI-Heuristic			CMI-Permutation			$\Delta MI-\chi^2$ [19]			MI-Permutation [14]			“Optimal”	
	#F	acc.	rank	#F	acc.	rank	#F	acc.	rank	#F	acc.	rank	#F	acc.
waveform (21)	17	84.2 ± 1.9	1	4	78.2 ± 2.1	2	3	76.1 ± 2.0	4	4	78.2 ± 2.1	2	11	84.7 ± 1.8
breast (30)	1	90.9 ± 1.5	2	1	90.9 ± 1.5	2	2	95.2 ± 1.2	1	1	90.9 ± 1.5	2	2	95.2 ± 1.2
heart (13)	6	82.6 ± 3.0	1	1	55.0 ± 4.4	3	2	76.9 ± 3.8	2	1	55.0 ± 4.4	3	6	82.6 ± 3.0
spect (22)	17	80.2 ± 4.2	1	2	78.2 ± 3.2	3	1	79.6 ± 3.2	4	3	79.1 ± 3.2	2	11	81.4 ± 4.2
ionosphere (34)	18	84.7 ± 2.9	1	8	81.9 ± 3.5	2	1	76.4 ± 3.9	4	8	81.9 ± 3.5	2	33	85.3 ± 3.0
parkinsons (22)	13	83.7 ± 3.7	1	4	84.6 ± 3.2	2	1	77.0 ± 5.0	4	3	84.6 ± 3.4	3	9	86.0 ± 3.9
semeion (256)	53	85.5 ± 1.1	1	40	82.7 ± 1.4	2	4	49.6 ± 1.7	4	40	82.7 ± 1.4	2	73	93.3 ± 1.3
Lung (325)	9	76.1 ± 7.9	2	5	74.2 ± 7.7	3	1	46.5 ± 7.5	4	11	73.0 ± 7.1	1	41	84.3 ± 6.5
Lymph (4026)	8	86.0 ± 5.3	3	65	89.6 ± 5.5	1	2	62.8 ± 6.5	4	64	89.6 ± 5.5	2	70	90.7 ± 5.4
Madelon (500)	3	69.5 ± 1.6	2	1	52.4 ± 1.7	3	4	76.7 ± 1.5	1	1	52.4 ± 1.7	3	4	76.7 ± 1.5
average rank			1.5			2.3			3.2			2.2		

As can be seen, $\Delta MI-\chi^2$ is likely to severely underestimate the number of features, accompanied by the lowest bootstrap accuracy. One possible reason is that $I(x^{cand}; y|S')$ does not precisely fit a

χ^2 distribution if the MB condition is not satisfied. CMI-permutation and MI-permutation always have the same ranks. This is because $I(S';y) + I(\{S - S'\};y|S') = I(S;y)$, a fixed value. Thus, it is equivalent to monitor the increment of $I(S';y)$ or the decrement of $I(\{S - S'\};y|S')$. On the other hand, it is surprising to find that CMI-heuristic performs the best in most datasets. This indicates that although the permutation test is effective to test the MB condition, $\varepsilon = 10^{-4}$ is a reliable threshold to speed up this test, as the permutation test is always time-consuming. If we look deeper, it is also interesting to see that different values of α in Rényi's entropy affect the results. This is to be expected, because the α is associated with the norm selected in the simplex to find the distance of the PDF to the center of the space. Effectively, different α weights change the effects of the tails, as explained in Reference [26]: Values of α smaller than 1 emphasize the importance of samples in the tails, values larger than 2 emphasize the modes of the distributions, whereas values close to 2 give equal weights to every sample, no matter where it is located in the distribution. Since classification uses a counting norm, values of α lower than 2 should be preferred. This is also the reason CMI-permutation and MI-permutation perform better with $\alpha = 1.01$. Finally, the Wilcoxon rank-sum test at 0.1 significance level, shown in Table 3, corroborates our analysis that our criteria perform equally to or slightly better than MI-permutation, but significantly better than $\Delta MI-\chi^2$.

Table 3. Summary of p -values and decisions (in parentheses) of Wilcoxon rank-sum test at 0.1 significance level on ranks of our criteria against $\Delta MI-\chi^2$ and MI-permutation. A p -value smaller than 0.1 indicates rejection of the null hypothesis that two criteria perform equally.

$\alpha = 1.01$		
	$\Delta MI-\chi^2$	MI-Permutation
CMI-heuristic	0.0781 (1)	0.5455 (0)
CMI-permutation	0.0561 (1)	0.9036 (0)
$\alpha = 2$		
	$\Delta MI-\chi^2$	MI-Permutation
CMI-heuristic	0.0081 (1)	0.0341 (1)
CMI-permutation	0.0587 (1)	0.7340 (0)

4. Conclusions

This paper suggests two simple stopping criteria, namely, CMI-heuristic and CMI-permutation, for information theoretic feature selection by monitoring the value of CMI estimated with the novel multivariate matrix-based Rényi's α -entropy functional. Experiments on 10 benchmark datasets indicate that (1) CMI is a more tractable quantity than MI, and (2) the non-parametric test (like the permutation test) is more effective than the parametric test with a prior distribution (like the χ^2 distribution), to guide early stopping in feature selection. Moreover, as an alternative to a permutation test, we also demonstrate a tiny threshold is sufficient to test the MB condition, which is very valuable for practitioners across different domains.

In the future, we will investigate the performance of our stopping criteria on big data. This is because with the tremendous growth of dataset sizes and dimensionality, the scalability of most feature selection methods could be jeopardized [33]. This problem is particularly important for our criteria, as the large-scale eigenvalue decomposition is always a challenging problem mathematically [34]. Moreover, we are also interested in testing our criteria on different types of data, including the linked data that is pervasive in social media and bioinformatics [33].

Author Contributions: methodology, S.Y. and J.C.P.; software, S.Y.; investigation, S.Y. and J.C.P.; writing—original draft preparation, S.Y.; writing—review and editing, S.Y. and J.C.P.; supervision, J.C.P.; funding acquisition, J.C.P.

Funding: This research was funded by the U.S. Office of Naval Research under grant number N00014-15-1-2103 and N00014-18-1-2306.

Acknowledgments: The authors thank the reviewers for their very helpful suggestions, which led to substantial improvements of the paper.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

MI	Mutual information
CMI	Conditional mutual information
RKHS	Reproducing kernel Hilbert space
MB	Markov blanket
PDF	Probability density function

References

1. Vergara, J.R.; Estévez, P.A. A review of feature selection methods based on mutual information. *Neural Comput. Appl.* **2014**, *24*, 175–186. [[CrossRef](#)]
2. Guyon, I.; Elisseeff, A. An introduction to variable and feature selection. *J. Mach. Learn. Res.* **2003**, *3*, 1157–1182.
3. Lu, Y.; Fan, Y.; Lv, J.; Noble, W.S. DeepPINK: Reproducible feature selection in deep neural networks. In Proceedings of the 2018 Conference on Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018; pp. 8690–8700.
4. Battiti, R. Using mutual information for selecting features in supervised neural net learning. *IEEE Trans. Neural Netw.* **1994**, *5*, 537–550. [[CrossRef](#)] [[PubMed](#)]
5. Meyer, P.E.; Schretter, C.; Bontempi, G. Information-theoretic feature selection in microarray data using variable complementarity. *IEEE J. Sel. Top. Signal Process.* **2008**, *2*, 261–274. [[CrossRef](#)]
6. Gurban, M.; Thiran, J.P. Information theoretic feature extraction for audio-visual speech recognition. *IEEE Trans. Signal Process.* **2009**, *57*, 4765–4776. [[CrossRef](#)]
7. Eriksson, T.; Kim, S.; Kang, H.G.; Lee, C. An information-theoretic perspective on feature selection in speaker recognition. *IEEE Signal Process. Lett.* **2005**, *12*, 500–503. [[CrossRef](#)]
8. Brown, G.; Pocock, A.; Zhao, M.J.; Luján, M. Conditional likelihood maximisation: a unifying framework for information theoretic feature selection. *J. Mach. Learn. Res.* **2012**, *13*, 27–66.
9. Fleuret, F. Fast binary feature selection with conditional mutual information. *J. Mach. Learn. Res.* **2004**, *5*, 1531–1555.
10. Ross, B.C. Mutual information between discrete and continuous data sets. *PLoS ONE* **2014**, *9*, e87357. [[CrossRef](#)]
11. Singha, S.; Shenoy, P.P. An adaptive heuristic for feature selection based on complementarity. *Mach. Learn.* **2018**, *107*, 1–45. [[CrossRef](#)]
12. Vinh, N.X.; Zhou, S.; Chan, J.; Bailey, J. Can high-order dependencies improve mutual information based feature selection? *Pattern Recognit.* **2016**, *53*, 46–58. [[CrossRef](#)]
13. Yu, S.; Giraldo, L.G.S.; Jenssen, R.; Principe, J.C. Multivariate Extension of Matrix-based Renyi's α -order Entropy Functional. *arXiv* **2018**, arXiv:1808.07912.
14. François, D.; Rossi, F.; Wertz, V.; Verleysen, M. Resampling methods for parameter-free and robust feature selection with mutual information. *Neurocomputing* **2007**, *70*, 1276–1288. [[CrossRef](#)]
15. Gómez-Verdejo, V.; Verleysen, M.; Fleury, J. Information-theoretic feature selection for the classification of hysteresis curves. In *International Work-Conference on Artificial Neural Networks*; Springer: Berlin/Heidelberg, Germany, 2007; pp. 522–529.
16. Koller, D.; Sahami, M. *Toward Optimal Feature Selection*; The Stanford University InfoLab: Stanford, CA, USA, 1996.
17. Yaramakala, S.; Margaritis, D. Speculative Markov blanket discovery for optimal feature selection. In Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM'05), Houston, TX, USA, 27–30 November 2005; pp. 809–812.
18. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 2012.

19. Vinh, N.X.; Chan, J.; Bailey, J. Reconsidering Mutual Information Based Feature Selection: A Statistical Significance View. In Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI'14), Quebec City, QC, Canada, 27–31 July 2014; pp. 2092–2098.
20. Good, P. *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses*; Springer: Berlin/Heidelberg, German, 2013.
21. Kraskov, A.; Stögbauer, H.; Grassberger, P. Estimating mutual information. *Phys. Rev. E* **2004**, *69*, 066138. [[CrossRef](#)] [[PubMed](#)]
22. Campos, L.M. A scoring function for learning Bayesian networks based on mutual information and conditional independence tests. *J. Mach. Learn. Res.* **2006**, *7*, 2149–2187.
23. Giraldo, L.G.S.; Rao, M.; Principe, J.C. Measures of entropy from data using infinitely divisible kernels. *IEEE Trans. Inf. Theory* **2015**, *61*, 535–548. [[CrossRef](#)]
24. Peng, H.; Long, F.; Ding, C. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Anal. Mach. Intell.* **2005**, *27*, 1226–1238. [[CrossRef](#)] [[PubMed](#)]
25. Rényi, A. *On Measures of Entropy and Information*; Hungarian Academy of Sciences: Budapest, Hungary, 1961.
26. Principe, J.C. *Information Theoretic Learning: Renyi's Entropy and Kernel Perspectives*; Springer: Berlin/Heidelberg, German, 2010.
27. Ohya, M.; Petz, D. *Quantum Entropy and Its Use*; Springer: Berlin/Heidelberg, German, 2004.
28. Bhatia, R. Infinitely divisible matrices. *Am. Math. Mon.* **2006**, *113*, 221–235. [[CrossRef](#)]
29. Müller-Lennert, M.; Dupuis, F.; Szehr, O.; Fehr, S.; Tomamichel, M. On quantum Rényi entropies: A new generalization and some properties. *J. Math. Phys.* **2013**, *54*, 122203. [[CrossRef](#)]
30. Mosonyi, M.; Hiai, F. On the Quantum Rényi Relative Entropies and Related Capacity Formulas. *IEEE Trans. Inf. Theory* **2011**, *57*, 2474–2487. [[CrossRef](#)]
31. Verdú, S. α -mutual information. In Proceedings of the 2015 IEEE Information Theory and Applications Workshop (ITA), San Diego, CA, USA, 1–6 February 2015; pp. 1–6.
32. Chang, C.C.; Lin, C.J. LIBSVM: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.* **2011**, *2*, 1–27. [[CrossRef](#)]
33. Li, J.; Liu, H. Challenges of feature selection for big data analytics. *IEEE Intell. Syst.* **2017**, *32*, 9–15. [[CrossRef](#)]
34. Martin, C.D.; Porter, M.A. The extraordinary SVD. *Am. Math. Mon.* **2012**, *119*, 838–851. [[CrossRef](#)]



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).