

# Data-Driven Representations for Testing Independence: Modeling, Analysis and Connection With Mutual Information Estimation

Mauricio E. Gonzalez, Jorge F. Silva <sup>✉</sup>, *Senior Member, IEEE*, Miguel Videla, and Marcos E. Orchard <sup>✉</sup>, *Member, IEEE*

**Abstract**—This work addresses testing the independence of two continuous and finite-dimensional random variables from the design of a data-driven partition. The empirical log-likelihood statistic is adopted to approximate the sufficient statistics of an oracle test against independence (that knows the two hypotheses). It is shown that approximating the sufficient statistics of the oracle test offers a learning criterion for designing a data-driven partition that connects with the problem of mutual information estimation. Applying these ideas in the context of a data-dependent tree-structured partition (TSP), we derive conditions on the TSP's parameters to achieve a strongly consistent distribution-free test of independence over the family of probabilities equipped with a density. Complementing this result, we present finite-length results that show our TSP scheme's capacity to detect the scenario of independence structurally with the data-driven partition as well as new sampling complexity bounds for this detection. Finally, some experimental analyses provide evidence regarding our scheme's advantage for testing independence compared with some strategies that do not use data-driven representations.

**Index Terms**—Independence testing, non-parametric learning, learning representations, data-driven partitions, tree-structure partitions, mutual information, consistency, finite-length analysis.

## I. INTRODUCTION

**T**HE problem of detecting independence from i.i.d. samples in a learning setting (distribution-free) is fundamental in statistics and has found numerous applications in statistical signal processing, data analysis, machine learning, inference, and decision problems [1]–[4]. For instance, the capacity to detect independent and conditional independent structures from data is relevant when having invariances and probabilistic symmetries in decision and machine learning models [5]–[8]. This capacity has been used in blind source separation, independent

component analysis (ICA) [1]–[4], [9] and the detection of associations in data [10]. Detecting independence from data has been systematically studied, and the literature is vast. Different non-parametric strategies have been proposed for this task using kernel-based statistics [11]–[14], distance-based approaches [15]–[19], histogram-based approaches [10], [20]–[24], log-likelihood statistics [20], [25], correlation-based approaches [26], [27],  $\phi$ -divergence estimators [20], [25], entropy and mutual information estimators [10], [10], [28]–[30], among many others.

A natural strategy, and the one we follow in this paper, is to partition the observation space (the binning approach) to evaluate the discrepancy (in some sense) between two empirical distributions [12], [15], [20]–[22], [29]. We highlight the following works using this approach: Gretton and Györfi [20] introduced a family of product data-independent partitions using the  $L_1$ -statistics [31] and the I-divergence statistics [32]. With these two approaches, sufficient conditions are established on the partition and other learning parameters to achieve strong consistency for detecting independence distribution-free in a continuous multivariate scenario. Székely *et al.* [15] also use partitions on the sample space, utilizing empirical distance correlation and distance covariance statistics to introduce a novel test of independence. The test is consistent and shows better empirical results than classical likelihood-ratio tests in scenarios with non-linear dependencies. Shi *et al.* [33] address the multivariate scenario by incorporating the concept of transportation-based ranks and signs. They propose plugging the center-outward distribution function into the mentioned distance covariance statistic [15] to determine a distribution-free test of independence.

More recently, Zhang [22] proposed a non-parametric test of independence assuming a rich class of models with known marginal distributions (uniform over  $[0, 1]$ ). This work proposed a binary expansion filtration of the space and deeply studied some symmetric properties to create a novel test that is uniformly consistent (on the power loss) and minimax for the power (in the sample size) assuming a large class of alternative distributions. Zhang *et al.* [34] presented a novel extension of the binary expansion filtration strategy [22] to a multidimensional setting by cleverly approximating an oracle *Neyman-Pearson* (NP) statistic. They propose a data-adaptive weight strategy (for the binary expansion) to approximate the ideal weight of the NP test. This new test unifies several important tests ( $\chi^2$  test

Manuscript received April 14, 2021; revised October 25, 2021; accepted November 29, 2021. Date of publication December 15, 2021; date of current version January 3, 2022. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Yuejie Chi. This work was supported in part by the CONICYT-Chile, Fondecyt 1210315, ANID-PFCHA/MagsterNacional/2019-22191445, and in part by the Advanced Center for Electrical and Electronic Engineering, Basal Project FB0008.

The authors are with the Information and Decision System Group, Department of Electrical Engineering, University of Chile, Santiago 412-3, Chile (e-mail: mauricio.gonzalez.gu@ing.uchile.cl; josilva@ing.uchile.cl; miguel.videla@ug.uchile.cl; morchard@ing.uchile.cl).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TSP.2021.3135689>, provided by the authors.

Digital Object Identifier 10.1109/TSP.2021.3135689

statistics, distance correlation) and shows promising empirical results.

On the important connection between testing independence and information measures [35], [36], we highlight the work of Kontoyiannis and Skoulariidou [30], where the problem of estimating directed information (directed information rate) between two discrete (finite alphabet) processes was explored based on a plug-in estimator. Importantly, the authors used this plug-in estimator to test the presence or absence of causality between the processes. Along the same lines, Suzuki [29] explored different strategies to estimate mutual information between discrete and continuous random variables and then evaluated those strategies experimentally for the task of testing independence. Finally, Reshef *et al.* [10] introduced the maximal information coefficient (MIC) to measure complex functional dependencies between a pair of variables by maximizing the (empirical) mutual information computed for collections of (adaptive) partitions of different resolutions.

In many of the methods mentioned that use partitions (or binning) for testing independence, a component that has not been explored systematically is the role played by data-driven partitions [10], [37], [38]. Data-driven partitions use data for the binning process (vector quantization) to construct the cells and define the final structure of the partition. This flexibility offers the capacity to better address inference and learning problems. Supporting this idea, data-driven representations have shown great approximation properties and better decision-estimation performance in numerous non-parametric learning problems [37]–[41].

In our context, the motivation that piqued our interest in data-driven partitions is the fact that under the hypothesis of independence, the trivial partition (the partition with one cell that contains all the space) is a sufficient representation of the problem. Our initial conjecture is that a well-designed adaptive data-driven partition could have the ability to detect (learn from data) this trivial solution by implementing a form of explicit regularization in its design. This adaptive ability could provide better detection of independence than non-adaptive partition schemes, which has prompted the exploration of data-driven methods for testing independence in this work.

Motivated by this, we look at the problem of designing a new learning criterion that selects a data-driven partition of the space (adaptive binning) as the optimal trade-off between estimation and approximation errors in a learning problem. To formulate this representation learning task, we adopt ideas from universal source coding to introduce a regret term that measures the discrepancy attributed to the use of empirical log-likelihood statistics — restricted over a partition — with respect to the oracle sufficient statistics of an ideal test that knows the true (two) probabilities: the oracle test against independence [36], [42]. Using this regret analysis, we establish a novel connection with the problem of mutual information (MI) estimation [30]. Furthermore, general conditions are derived to obtain a strongly consistent test of independence in the strong (almost-sure and distribution-free) sense introduced in [20].

We apply this framework in the context of tree-structured partitions (TSP) [39], which is the main application focus of this

work. TSP algorithms were selected for the following important reasons: Their implementation is simple, TSP have a binary structure that can be used to address learning tasks efficiently for a large class of optimization problems (minimum cost-tree pruning) [43]–[45], and they have been shown to be expressive (with good approximation properties) when compared with other partition strategies in non-parametric learning and decision tasks [37], [37]–[41], [43], [44], [46].

On the proposed test, we derive concrete connections with results on mutual information estimation [39]. From these connections, we established new design conditions to obtain a strongly consistent test of independence and, more importantly, non-asymptotic results that express our framework's capacity to approximate the information term of the oracle (ideal) test with a finite sample-size. Going a step further in this finite length performance analysis, we study our scheme's capacity to detect independence structurally with the underlying data-driven partition, which was one of the original motivations used to explore data-driven partition for this task.

Indeed, we show that under the independence assumption, our data-driven partition collapses to the trivial solution with one cell (implying that the resulting MI estimator is zero) with a finite sample-size almost surely. This is a remarkable property attributed to our scheme's capacity to adapt its representation to the sufficient statistics of the problem, which is the trivial partition in the context of independence. From this ability, we improve our original result concerning our test's consistency (density-free) and provide refined sampling complexity bounds for detecting the scenario of independence.

To implement our approach, we propose a learning algorithm that offers a computationally efficient implementation of our data-driven strategy for testing independence. The algorithm is divided in two important phases (a growing and pruning stage), which are both shown to have a polynomial complexity on the side of the problem (the sample size). Finally, we provide empirical evidence of the advantage of our strategy in some controlled simulation scenarios. For this last analysis, we introduce concrete non-asymptotic metrics (in the form of sampling complexity indicators) to measure the test's ability to detect the correct hypothesis with a finite number of samples.

A preliminary version of this work was presented in [47]. This work significantly expands the technical contribution in [47] and explores implementations and experimental analyses.

### A. Organization of the Paper

Section II introduces the problem and basic notations. Section III presents the general histogram-based scheme, and Section IV introduces the regret-based analysis proposed in this work and its connection with the problem of MI estimation. Section IV-B presents a general consistency result for our problem (Theorem 1). Section V introduces the family of tree-structured data-driven partitions and presents consistency results (Theorem 2, Lemma 2 and Corollary 1). Section VI analyzes some finite-length properties (Theorem 2) and elaborates on the structural capacity of our scheme to detect independence under the null hypothesis (Lemma 3, and Theorems 3 and 4). Sections VII and

VIII cover implementation and empirical analysis, respectively. Finally, Section IX concludes with some discussion and future directions. The proofs are presented in the Appendix.

## II. PRELIMINARIES

Let us consider two random variables  $X$  and  $Y$  taking values in  $\mathbb{R}^p$  and  $\mathbb{R}^q$ , respectively, with joint distribution  $P$  in  $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ , where  $d = p + q$  and  $\mathcal{B}(\mathbb{R}^d)$  is the Borel sigma field. Let us denote by  $\mathcal{P}_0$  the class of probabilities in  $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$  under which  $(X, Y)$  are independent, meaning that if  $P \in \mathcal{P}_0$ , then  $P(A \times B) = P(\mathbb{R}^p \times B) \cdot P(A \times \mathbb{R}^q)$  for any  $A \in \mathcal{B}(\mathbb{R}^p)$  and  $B \in \mathcal{B}(\mathbb{R}^q)$ .

Given  $n$  i.i.d. samples of  $(X, Y)$ , denoted by  $Z_1^n \equiv (Z_1 \equiv (X_1, Y_1), Z_2 \equiv (X_2, Y_2), \dots, Z_n \equiv (X_n, Y_n)) \in \mathbb{R}^d$  driven by the model  $P \in \mathcal{P}(\mathbb{R}^d)$ , the problem is to decide from  $Z_1^n$  whether  $X$  and  $Y$  are independent (meaning that  $P \in \mathcal{P}_0$ ) or, alternatively,  $X$  and  $Y$  have some statistical dependency, i.e.,  $P \notin \mathcal{P}_0$ . In this context, a decision rule  $\phi_n(\cdot)$  of length  $n$  is a function from  $(\mathbb{R}^d)^n$  to  $\{0, 1\}$ , where  $\phi_n(z_1^n) = 0$  means that the rule decides that the underlying probability producing  $z_1^n$  belongs to  $\mathcal{P}_0$ . Then, for any decision rule  $\phi_n \in \Pi_n^1$  of length  $n$ , we recognize two errors:

- Assuming that  $P \in \mathcal{P}_0$  ( $\mathcal{H}_0$ ), the non-detection of independence between  $X$  and  $Y$  is measured by<sup>2</sup>

$$\mathbb{P}(\phi_n(Z_1^n) = 1) = P^n(\{z_1^n : \phi_n(z_1^n) = 1\}),$$

- Assuming that  $P \notin \mathcal{P}_0$  ( $\mathcal{H}_1$ ), the false detection of independence between  $X$  and  $Y$  is measured by

$$\mathbb{P}(\phi_n(Z_1^n) = 0) = P^n(\{z_1^n : \phi_n(z_1^n) = 0\}).$$

The following is the classical notion of strong-consistency used to evaluate the asymptotic goodness of a universal scheme (collection of rules of different lengths) for detecting independence from data [20].

*Definition 1:* Given a scheme  $\xi = \{\phi_n \in \Pi_n, n \geq 1\}$ , where  $\phi_n$  is a decision rule of length  $n$ , we say that  $\xi$  is strongly consistent for detecting independence if

- under  $\mathcal{H}_0$ :  $(\phi_n(Z_1^n))_{n \geq 1}$  reaches 0 eventually with probability one, or  $\mathbb{P}$ -almost surely (a.s.),
- under  $\mathcal{H}_1$ :  $(\phi_n(Z_1^n))_{n \geq 1}$  reaches 1 eventually  $\mathbb{P}$ -a.s.

### A. The Divergence and Mutual Information

Let  $P, Q$  be two probability measures in  $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$  such that  $P \ll Q$ , and let us consider  $\pi = \{A_i, i \in \mathcal{I}\}$  a measurable partition of  $\mathbb{R}^d$  where  $\mathcal{I}$  is finite or countable. The divergence of  $P$  with respect to  $Q$  restricted over  $\pi$  (or the sub-sigma field induced by  $\pi$  denoted by  $\sigma(\pi)$ ) is given by [32], [36]

$$D_{\sigma(\pi)}(P||Q) \equiv \sum_{i \in \mathcal{I}} P(A_i) \log \frac{P(A_i)}{Q(A_i)}. \quad (1)$$

<sup>1</sup> $\Pi_n$  denotes the collection of binary rules acting on  $\mathbb{R}^{dn}$ .

<sup>2</sup> $\mathbb{P}$  denotes the entire process distribution of  $(Z_n)_{n \geq 1}$  and  $P^n$  is the  $n$ -fold probability in  $\mathbb{R}^{dn}$  induced by the marginal  $P$ .

The divergence of  $P$  with respect to  $Q$  is [32], [36]

$$D(P||Q) \equiv \sup_{\pi \in \mathcal{Q}(\mathbb{R}^d)} D_{\sigma(\pi)}(P||Q) \quad (2)$$

where  $\mathcal{Q}(\mathbb{R}^d)$  is the collection of finite measurable<sup>3</sup> partitions of  $\mathbb{R}^d$ .

## III. THE EMPIRICAL LOG-LIKELIHOOD STATISTICS FROM A DATA-DRIVEN PARTITION

In this work, we adopt a histogram-based log-likelihood statistics approach [20]. We interpret this approach as an empirical version of the Neyman-Pearson (NP) test against independence [42]. The ideal (oracle) NP test against independence decides whether the samples belong to the following two known cases:  $P$  (for some  $P \notin \mathcal{P}_0$ ) or  $Q^*(P) \in \mathcal{P}_0$ , where  $Q^*(A \times B) = P(\mathbb{R}^p \times B) \cdot P(A \times \mathbb{R}^q)$  for any  $A \in \mathcal{B}(\mathbb{R}^p)$  and  $B \in \mathcal{B}(\mathbb{R}^q)$ .<sup>4</sup>

The specific approach proposed in this work is a two-stage process that uses the data  $Z_1^n \sim P^n$  twice: first, to estimate  $P$ , and its projected version  $Q^*(P)$  (over  $\mathcal{P}_0$ ), with the caveat that the estimation is restricted over the events of a finite partition of the joint space  $\mathbb{R}^d$  (i.e., a histogram-based estimation of  $P$  and  $Q^*(P)$ , respectively); second, to compute an empirical version of the log likelihood-ratio to decide (using a threshold) if  $\hat{P} = \hat{Q}^*$  (independence) or  $\hat{P} \neq \hat{Q}^*$  (non-independence).

There are two elements of learning design that determine our empirical test. The first is a data-driven partition rule denoted by  $\pi_n(\cdot)$ , which is a function that maps sequences in  $(\mathbb{R}^d)^n$  into a finite measurable partitions of  $\mathbb{R}^d$ . The second element is a non-negative threshold denoted by  $a_n \in \mathbb{R}^+$ . Then given a pair  $(\pi_n(\cdot), a_n)$  and some data  $z_1^n \in \mathbb{R}^{dn}$ , the data-driven test is constructed as follows:

- 1) Use the quantization  $\pi_n(z_1^n) = \{A_i, i = 1, \dots, |\pi_n(z_1^n)|\} \subset \mathcal{B}(\mathbb{R}^d)$  to estimate  $P$  over the cells  $\pi_n(z_1^n)$ :

$$\hat{P}_n(A) \equiv \frac{1}{n} \sum_{i=1}^n \mathbf{1}_A(z_i) \text{ and} \quad (3)$$

$$\hat{Q}_n^*(\underbrace{A^1 \times A^2}_{A=}) \equiv \hat{P}_n(A^1 \times \mathbb{R}^q) \cdot \hat{P}_n(\mathbb{R}^p \times A^2), \quad (4)$$

for any  $A \in \pi_n(z_1^n)$ . In (4) we assume that the cells of  $\pi_n(z_1^n)$  have a product structure, i.e.,  $A = A^1 \times A^2$  where  $A^1 \in \mathcal{B}(\mathbb{R}^p)$  and  $A^2 \in \mathcal{B}(\mathbb{R}^q)$ .<sup>5</sup>

- 2) Project the data  $z_1^n$  over the cells of  $\pi_n(z_1^n)$ . A simple projection (or representation) function is the following:

$$O_{\pi_n}(z) \equiv \sum_{j=1}^{|\pi_n(z_1^n)|} j \cdot \mathbf{1}_{A_j}(z) \in \{1, \dots, |\pi_n(z_1^n)|\}. \quad (5)$$

The projected data is given and denoted by  $o_1^n \equiv (o_1 = O_{\pi_n}(z_1), \dots, o_n = O_{\pi_n}(z_n))$ .

- 3) Compute the log-likelihood ratio of the empirical distributions in (3) and (4) using the quantized (or projected over

<sup>3</sup> $\mathcal{B}(\mathbb{R}^d)$  is the sigma field of Borel sets.

<sup>4</sup> $Q^*(P)$  can be interpreted as the projection of  $P$  on  $\mathcal{P}_0$  in the information divergence sense, i.e.,  $Q^*(P)$  is the solution of  $\min_{Q \in \mathcal{P}_0} D(P||Q)$ .

<sup>5</sup>This event-wise product structure on the cells of  $\pi_n(z_1^n)$  is needed to estimate both  $P$  and  $Q^*$  from i.i.d. samples of  $P$ .

$\sigma(\pi_n(z_1^n))$  data  $o_1^n$ . In particular, we consider the following log-likelihood ratio per sample as

$$\hat{i}_{\pi_n}(o_1^n) \equiv \frac{1}{n} \log \frac{\hat{P}_{O_{\pi_n}}(o_1) \hat{P}_{O_{\pi_n}}(o_2) \cdots \hat{P}_{O_{\pi_n}}(o_n)}{\hat{Q}_{O_{\pi_n}}(o_1) \hat{Q}_{O_{\pi_n}}(o_2) \cdots \hat{Q}_{O_{\pi_n}}(o_n)}, \quad (6)$$

where  $\hat{P}_{O_{\pi_n}}$  and  $\hat{Q}_{O_{\pi_n}}^*$  denote the empirical distributions of the quantized random variable  $O_{\pi_n}$  in its representation space  $\{1, \dots, |\pi_n(z_1^n)|\}$ , meaning that by construction  $\hat{P}_{O_{\pi_n}}(j) = \hat{P}_n(A_j)$  and  $\hat{Q}_{O_{\pi_n}}(j) = \hat{Q}_n^*(A_j)$  for all  $j \in \{1, \dots, |\pi_n(z_1^n)|\}$ .

4) Finally, the decision is given by the following rule:

$$\phi_{\pi_n, a_n}(o_1^n) \equiv \begin{cases} 0 & \text{if } \hat{i}_{\pi_n}(o_1^n) < a_n \\ 1 & \text{if } \hat{i}_{\pi_n}(o_1^n) \geq a_n \end{cases}. \quad (7)$$

The second step introduced  $O_{\pi_n}(Z)$  in (5). This object captures the role played by the data-driven partition ( $\pi_n$ ) in our log-likelihood statistics in (6).

#### IV. REGRET ANALYSIS: THE REPRESENTATION LEARNING PROBLEM

To analyze the quality of the proposed test, in this section we compare the proposed statistics in (6) with the statistics used by an oracle NP test. We consider the data-driven partition in the pair  $(\pi_n(\cdot), a_n)$  as a learning agent. For this analysis, let us first fix a finite partition  $\pi = \{A_i, i = 1, \dots, \{1, \dots, J\}\}$ . Given i.i.d. realizations  $Z_1, \dots, Z_n$  from  $P \in \mathcal{P}(\mathbb{R}^d)$ , we consider the empirical-quantized log-likelihood statistics in (6). Adopting the notion of regret from universal source coding [35], let us consider as a reference the true likelihood ratio of  $P$  with respect to  $Q^*(P)$  associated with the expression in (6) but with no quantization effects, i.e.,<sup>6</sup>

$$i_n(z_1, \dots, z_n) \equiv \frac{1}{n} \log \frac{dP}{dQ^*}(z_1) \cdot \frac{dP}{dQ^*}(z_2) \cdots \frac{dP}{dQ^*}(z_n). \quad (8)$$

$i_n(z_1, \dots, z_n)$  is the ideal (oracle) information term used by the NP test against independence [35]. We measure the (sample-wise) regret of  $\pi$  as

$$\hat{i}_n(z_1, \dots, z_n) - \hat{i}_\pi(O_\pi(z_1), \dots, O_\pi(z_n)), \quad (9)$$

which measures the discrepancy between the empirical-quantized statistics in (6) and the oracle term in (8).

##### A. Connection With Divergence Estimation

The regret in (9) has two error sources: one associated with the quantization of the space (the representation quality of  $\pi$ ) and the other associated with the fact that we use the empirical distribution  $\hat{P}_n$  in (3) instead of the true model  $P$ . To isolate these components, it is useful to introduce *the oracle-quantized information* given by

$$i_\pi(o_1^n) \equiv \frac{1}{n} \log \prod_{i=1}^n \frac{P_{O_\pi}(o_i)}{Q_{O_\pi}^*(o_i)} \text{ with } o_1^n \in \{1, \dots, J\}^n, \quad (10)$$

<sup>6</sup>By construction  $P \ll Q^*$ ; therefore, the RN derivative  $\frac{dP}{dQ^*}$  is well defined in (8) where it is clear that if  $(\log \frac{dP}{dQ^*}(z))_{z \in \mathbb{R}^d} \in \ell_1(P)$ , then  $I(X, Y) = \mathbb{E}_{Z \sim P} \{\log \frac{dP}{dQ^*}(Z)\} < \infty$ .

where  $P_{O_\pi}(j) = P(A_j)$  and  $Q_{O_\pi}^*(j) = Q^*(A_j)$  for all  $j \in \{1, \dots, J\}$  are short-hand for the true probabilities induced by  $O_\pi(\cdot)$ ,  $P$  and  $Q^*$  in  $\{1, \dots, J\}$ . Then, the regret in (9) can be decomposed in two components:

$$\begin{aligned} & \hat{i}_n(z_1, \dots, z_n) - \hat{i}_\pi(O_\pi(z_1), \dots, O_\pi(z_n)) = \\ & \underbrace{\hat{i}_n(z_1, \dots, z_n) - \hat{i}_\pi(O_\pi(z_1), \dots, O_\pi(z_n))}_I + \\ & \underbrace{i_\pi(O_\pi(z_1), \dots, O_\pi(z_n)) - \hat{i}_\pi(O_\pi(z_1), \dots, O_\pi(z_n))}_{II}. \end{aligned} \quad (11)$$

The first term (I) on the right hand side (RHS) of Eq.(11) captures an information loss attributed to the quantization (5) (the approximation error). This expression converges as  $n$  tends to infinity to  $D(P||Q^*) - D_{\sigma(\pi)}(P||Q^*) \geq 0$  almost surely. The second term (II) on the RHS of (11) captures the discrepancy in the information density attributed to the use of empirical distributions. Crucially, these two error sources are driven by the partition  $\pi$ . At this point, it is worth noticing that the empirical information term can be expressed as  $\hat{i}_\pi(O_\pi(z_1), \dots, O_\pi(z_n)) = \sum_{A \in \pi} \hat{P}_n(A) \log \frac{\hat{P}_n(A)}{\hat{Q}_n^*(A)}$ , which is a histogram-based estimator of the divergence restricted over the cells of  $\pi$  [39], [40], [48]. From these results, the RHS of (11) can be re-structured as follows

$$\begin{aligned} & \hat{i}_n(Z_1, \dots, Z_n) - \hat{i}_\pi(O_\pi(Z_1), \dots, O_\pi(Z_n)) = \\ & \underbrace{\hat{i}_n(Z_1, \dots, Z_n) - D(P||Q^*)}_I + \underbrace{D(P||Q^*) - D_{\sigma(\pi)}(P||Q^*)}_{II} + \\ & \underbrace{D_{\sigma(\pi)}(P||Q^*) - D_{\sigma(\pi)}(\hat{P}_n||\hat{Q}_n^*)}_{III}. \end{aligned} \quad (12)$$

The first term (I) on the RHS of (12) goes to zero with probability one (by the strong law of large numbers assuming that  $D(P||Q^*) < \infty$ ) independent of  $\pi$ . Therefore, from the perspective of evaluating the effect of  $\pi$  in the regret, this term can be overlooked. On the other hand, the second term (II) captures the approximation error, or what we lost in discrimination (as the number of samples goes to infinity) when projecting the data over  $\sigma(\pi)$  with respect to the lossless term in (8). Finally, the last term (III) measures the estimation error, which is the discrepancy between the true (oracle) distributions and the empirical distributions in the information divergence sense over events of  $\sigma(\pi)$  (see Eq.(1)).

1) *The Representation Problem.* Here we introduce the main design problem of this work, which is to find a data-driven partition that offers an optimal balance between the two relevant terms presented in (12). For the estimation error, the idea is to adopt distribution-free (universal) error bounds for  $|D_{\sigma(\pi)}(P||Q^*) - D_{\sigma(\pi)}(\hat{P}_n||\hat{Q}_n^*)|$  of the form (more details in Section V-B):

$$\mathbb{P} \left( \left| D_{\sigma(\pi)}(P||Q^*) - D_{\sigma(\pi)}(\hat{P}_n||\hat{Q}_n^*) \right| \leq r_n(\pi) \right) \geq 1 - \delta,$$

where  $r_n(\pi)$  is the confidence interval for the confidence probability  $1 - \delta$ . Equipped with this result, we could use the following upper bound for the regret:  $D(P||Q^*)$

–  $\hat{i}_\pi(O_\pi(z_1), \dots, O_\pi(z_n)) \leq [D(P||Q^*) - D_{\sigma(\pi)}(P||Q^*)] + r_n(\pi)$ , with high probability. Finally, we could formulate the problem of selecting  $\pi$  over a family of representations  $\Pi$  (subset of  $\mathcal{Q}(\mathbb{R}^d)$ ) as the solution of the following *regularized info-max problem*:

$$\pi_n^* \equiv \arg \max_{\pi \in \Pi} D_{\sigma(\pi)}(P||Q^*) - r_n(\pi). \quad (13)$$

The learning task in (13) is still intractable. It resembles an oracle learner agent (teacher) that selects  $\pi$  in  $\Pi$  solving the trade-off between the two errors, one of which needs the true model  $P$ . Section V addresses a tractable info-max version of (13), where instead of  $P$  the empirical distribution  $\hat{P}$  in (3) is adopted.

### B. Strong Consistency: A Basic Requirement

In this subsection, we introduce the idea of measuring the discrepancy between  $\hat{i}_{\pi_n}(O_{\pi_n}(z_1), \dots, O_{\pi_n}(z_n))$  and the ideal (oracle) information  $i_n(z_1, \dots, z_n)$ . From this perspective, we could introduce a new notion of consistency based on this learning objective.

*Definition 2:* Our scheme  $\{(\pi_n, a_n), n \geq 1\}$  is said to be *strongly consistent on the regret* if for any  $P \in \mathcal{P}(\mathbb{R}^d)$  and i.i.d. process  $(Z_n)_{n \geq 1}$  with  $Z_i \sim P$ , it follows that

$$\lim_{n \rightarrow \infty} \left| i_n(Z_1, \dots, Z_n) - \hat{i}_{\pi_n}(O_{\pi_n}(Z_1), \dots, O_{\pi_n}(Z_n)) \right| = 0,$$

$\mathbb{P}$ -almost surely (a.s.).

A simple result follows:

*Proposition 1:* If  $\{(\pi_n, a_n), n \geq 1\}$  is strongly consistent on the regret, then  $\hat{i}_{\pi_n}(O_{\pi_n}(Z_1), \dots, O_{\pi_n}(Z_n))$  is a strongly consistent estimator of the MI between  $X$  and  $Y$ .<sup>7</sup>

Returning to our original problem, the next result offers sufficient conditions for a data-driven scheme  $\{(\pi_n, a_n), n \geq 1\}$  to be strongly consistent for detecting independence (Def. 1).

*Theorem 1:* Let  $\{(\pi_n, a_n), n \geq 1\}$  be the data-driven scheme of Section III.

- i) If  $(a_n)_n$  is  $o(1)$ ,
- ii)  $\{\pi_n, n \geq 1\}$  is strongly consistent on the regret (Def. 2),
- iii) under  $\mathcal{H}_0$ ,  $(\hat{i}_{\pi_n}(O_{\pi_n}(Z_1), \dots, O_{\pi_n}(Z_n)))_{n \geq 1}$  is  $o(a_n)$  in the sense that  $\lim_{n \rightarrow \infty} \frac{\hat{i}_{\pi_n}(O_{\pi_n}(Z_1), \dots, O_{\pi_n}(Z_n))}{a_n} = 0$   $\mathbb{P}$ -a.s.,

then  $(\phi_{\pi_n, a_n}(O_{\pi_n}(\cdot)))_{n \geq 1}$  in (7) is strongly consistent for detecting independence (Def. 1).

The proof is presented in Appendix A.

Consistency is a basic asymptotic requirement that is non-sufficient when looking into practical applications that operate with a finite sample size. Hence, in the next sections, we present a test that is strongly consistent (Th. 3) but also satisfies some relevant non-asymptotic properties in the form of finite-length performance results (Ths. 2, 4 and Lemma 3).

<sup>7</sup>From Definition 2, it follows directly that  $\lim_{n \rightarrow \infty} \hat{i}_{\pi_n}(O_{\pi_n}(Z_1), \dots, O_{\pi_n}(Z_n)) = D(P||Q^*) = I(X, Y)$ ,  $\mathbb{P}$ -almost surely.

## V. TREE-STRUCTURED DATA-DRIVEN PARTITIONS

In this section, an empirical version of (13) is studied considering for  $\Pi$  a dictionary of tree-structured data-driven partitions (TSP). This TSP family was introduced in [39] for the problem of MI estimation. The focus of this section is to demonstrate its potential for the problem of testing independence.

### A. The Collection of Partitions

The construction of this family begins by defining a “root” node to index the trivial partition  $A_{root} \equiv \{\mathbb{R}^d\}$ . The process continues by selecting a coordinate in  $\{1, \dots, d\}$  to project the data (1D projection), then order the projected data (scalar values), select the median of the ordered sequence, and finally create a statistically equivalent partition of  $A_{root}$  using the selected coordinate axis and the median to split the cell. Two new cells are created from  $A_{root}$  with almost half the sample points in each (exactly half when  $n$  is an even integer).

These new cells are indexed as the left and right children of the “root” denoted by  $(l(root), r(root))$ ; i.e., the new cells are  $A_{l(root)}$  and  $A_{r(root)}$  where  $A_{root} = A_{l(root)} \cup A_{r(root)}$ . The process continues by selecting a new coordinate in  $\{1, \dots, d\}$ , where the statistically equivalent binary splitting criterion is iterated in  $A_{l(root)}$  and  $A_{r(root)}$  until a stopping condition is met. The stopping criterion adopted here imposed a minimum empirical probability in each created cell that we denote by  $b_n \in (0, 1)$ . Therefore, at the end of this growing binary process, a collection of nodes is produced  $\mathcal{I} \equiv \{root, l(root), r(root), \dots\}$  associated with a binary tree (by construction) and the nodes’ respective cells  $\{A_v, v \in \mathcal{I}\}$  where  $\hat{P}_n(A_v) \geq b_n$  for any  $v \in \mathcal{I}$ .

Using the Breiman *et al.* [43] convention, a binary tree  $T$  is a collection of nodes in  $\mathcal{I}$ : one node of degree 2 (the “root”), and the remaining nodes of degree 3 (internal nodes) or degree 1 (leaf or terminal nodes) [43]. In this convention, the full-tree is denoted and given by  $T_{b_n}^{full} \equiv \mathcal{I}$ . Importantly, if  $\tilde{T} \subset T_{b_n}^{full}$  and  $\tilde{T}$  is a binary tree by itself, then we say that  $\tilde{T}$  is a subtree of  $T_{b_n}^{full}$ . Moreover, if both have the same root, we say that  $\tilde{T}$  is a *pruned version* of  $T_{b_n}^{full}$ , and we denote this by  $\tilde{T} \ll T_{b_n}^{full}$ . Finally, if we denote by  $\mathcal{L}(T)$  the leaf nodes of an arbitrary tree  $T \ll T_{b_n}^{full}$ , it is simple to verify that  $\pi_T \equiv \{A_v, v \in \mathcal{L}(T)\}$  is a data-driven partition of  $\mathbb{R}^d$  indexed by  $T$  where every cell in  $\pi_T$  has the desired product structure in (4) [39].<sup>8</sup>

### B. Regularized (Empirical) Information Maximization

To address the info-max design objective in (13) using our (data-driven) tree-indexed family  $\Pi_n \equiv \{\pi_T : T \ll T_{b_n}^{full}\}$ , we have to first find an expression for  $r_n(\pi)$  in (13) for any  $\pi_T \in \Pi_n$ . The next result offers the following:

*Lemma 1:* (Silva *et al.* [39, Th. 1]) Let  $\mathcal{G}_{b_n}^k \equiv \{T \ll T_{b_n}^{full} : |T| = k\}$  be the family of pruned TSPs of size  $k$ . Then,  $\forall k \in \{1, \dots, |T_{b_n}^{full}|\}$ ,  $\forall n > 0$ , and any small  $\delta \in (0, 1)$ , there is a

<sup>8</sup>The interested reader is referenced to [39] for a systematic explanation of this TSP construction.

threshold  $\epsilon_c(\cdot)$  where

$$\mathbb{P} \left( \sup_{T \in \mathcal{G}_{b_n}^k} \left| D_{\sigma(\pi_T)}(P||Q^*) - D_{\sigma(\pi_T)}(\hat{P}_n||\hat{Q}_n^*) \right| \leq \epsilon_c(\cdot) \right) \geq 1 - \delta, \quad (14)$$

where specifically  $\epsilon_c(n, b_n, d, \delta, k) =$

$$\frac{24\sqrt{2}}{b_n\sqrt{n}} \sqrt{\ln(8/\delta) + k[(d+1)\ln(2) + d\ln(n)]}. \quad (15)$$

Importantly the bound in (14) is distribution-free [39], and a function of  $b_n$  and  $k$  (the size of the family  $\mathcal{G}_{b_n}^k$ ). From this distribution-free concentration result, the union bound tells us that the following events in  $\mathbb{R}^{dn}$

$$\left\{ \sup_{T \in \mathcal{G}_{b_n}^k} \left| D_{\sigma(\pi_T)}(P||Q^*) - D_{\sigma(\pi_T)}(\hat{P}_n||\hat{Q}_n^*) \right| \leq r_{b_n, \delta}(k) \right\}, \quad (16)$$

with  $r_{b_n, \delta}(k) \equiv \epsilon_c(n, b_n, d, \delta \cdot b_n, k)$ , happen simultaneously for any  $k = 1, \dots, |T_{b_n}^{full}|$  with probability at least  $1 - \delta$  (with respect to  $\mathbb{P}$ ). Equipped with these penalizations, i.e.,  $r_{b_n, \delta}(|T|)$  in (16) for any  $T \ll T_{b_n}^{full}$ , the empirical version of the info-max problem in (13) is

$$\hat{T}_{b_n, \delta_n} \equiv \arg \max_{T \ll T_{b_n}^{full}} D_{\sigma(\pi_T)}(\hat{P}_n||\hat{Q}_n^*) - r_{b_n, \delta_n}(|T|). \quad (17)$$

Finally, it is important to note that both  $(b_n)_{n \geq 1}$  and  $(\delta_n)_{n \geq 1}$  determine the trees  $(\hat{T}_{b_n, \delta_n})_{n \geq 1}$  and TS partitions that we denote here by  $(\pi_{b_n, \delta_n})_{n \geq 1}$ . In addition, if we include the thresholds  $(a_n)_{n \geq 1}$ , we denote by  $(\phi_{b_n, \delta_n, a_n}(\cdot))_{n \geq 1}$  the rules induced by  $(\pi_{b_n, \delta_n})_{n \geq 1}$  and  $(a_n)_{n \geq 1}$  in (7).

### C. Consistency: A Preliminary Analysis

The following results show that the TSP scheme in (17) is strongly consistent on the regret (Def. 2).

*Lemma 2:* Let us assume that  $P$  has a density<sup>9</sup> in  $\mathbb{R}^d$ .

**i)** Under the conditions that  $(b_n) \approx (n^{-l})$  for  $l \in (0, 1/3)$ ,  $(\delta_n)$  is  $o(1)$  and  $(1/\delta_n)$  is  $\mathcal{O}(e^{n^{1/3}})$ , we have that

$$\lim_{n \rightarrow \infty} \left| i_n(Z_1, \dots, Z_n) - \hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n)) \right| = 0, \quad \mathbb{P}\text{-a.s.} \quad (18)$$

**ii)** Assuming that  $X$  and  $Y$  are independent ( $\mathcal{H}_0$ ), if we select  $(1/\delta_n)_n \approx (e^{n^{1/3}})$ , it follows that for any  $p > 0$

$$(\hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n)))_{n \geq 1} \text{ is } o(n^{-p}), \quad \mathbb{P}\text{-a.s.} \quad (19)$$

The proof is presented in Appendix B.

In addition to achieve consistency on the regret, stated in part i), the part ii) of this result shows that under  $\mathcal{H}_0$  the regret achieves a super polynomial velocity of convergence to zero: note that  $i_n(Z_1, \dots, Z_n) = 0$  with probability one under  $\mathcal{H}_0$ . Therefore, the empirical information rate

<sup>9</sup> $P$  is absolutely continuous with respect to the Lebesgue measure in  $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ .

$\hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n))$  tends to zero faster than any polynomial order in  $n$  with probability one, which is a remarkable capacity to detect this condition from data. In light of Theorem 1, we could have a range of admissible vanishing thresholds  $(a_n)_{n \geq 1}$  where strong consistency for detecting independence can be achieved (Def. 1). This is stated in the following result.

*Corollary 1:* Let us assume the setting and conditions on  $(b_n)_{n \geq 1}$  and  $(\delta_n)_{n \geq 1}$  stated in Lemma 2 part ii) for the TSP scheme. If  $(a_n)_{n \geq 1}$  is  $\mathcal{O}(n^{-q})$  for some arbitrary  $q > 0$ , then the TSP scheme  $\{\phi_{b_n, \delta_n, a_n}, n \geq 1\}$  is strongly consistent for detecting independence (Def. 1).

*Proof:* The proof derives directly from Theorem 1 and the two results obtained in Lemma 2.  $\square$

## VI. MAIN FINITE-LENGTH RESULTS

In this section, we focus on finite-length (non-asymptotic) results. We establish conditions where the solution of (17) nearly matches the performance of its equivalent oracle version in (13) with high probability. This non-asymptotic result is instrumental to show later one of the main findings of our work: the capacity of our scheme to detect  $\mathcal{H}_0$  with the structure of the partition.

*Theorem 2:* Under the conditions that  $(b_n) \approx (n^{-l})$  for  $l \in (0, 1/3)$ ,  $(\delta_n)$  is  $o(1)$  and  $(1/\delta_n)$  is  $\mathcal{O}(e^{n^{1/3}})$ , we have that

$$\underbrace{\hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n))}_{\text{zero regret regime}} = i_n(Z_1, \dots, Z_n) = 0, \quad (20)$$

holds with  $\mathbb{P}$ -probability  $1 - \epsilon$ .

**ii)** under  $\mathcal{H}_1$ : for any  $\epsilon > 0$  there is  $N(\epsilon)$  such that  $\forall n \geq N(\epsilon)$  the bound

$$\begin{aligned} & i_n(Z_1, \dots, Z_n) - \hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n)) \\ & \leq i_n(Z_1, \dots, Z_n) - I(X, Y) + \\ & \min_{T \ll T_{b_n}^{full}} [D(P||Q^*(P)) - D_{\sigma(\pi_T)}(P||Q^*(P))] + 2r_{b_n, \delta_n}(|T|). \end{aligned} \quad (21)$$

holds with  $\mathbb{P}$ -probability  $1 - \epsilon$ .

The proof of this result is presented in Appendix C.

This result presents two optimality bounds for the information term of the TSP scheme under the two main hypotheses ( $\mathcal{H}_0$  and  $\mathcal{H}_1$ ) of our problem. Under  $\mathcal{H}_0$ , our regularization approach is capable of detecting this structure (with an arbitrary high probability  $1 - \epsilon$ ) in the sense that  $\hat{T}_{b_n, \delta_n}$  in (17) reduces to the trivial partition  $\{\mathbb{R}^d\}$ . From this, we obtain the zero regret condition  $\hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n)) = i_n(Z_1, \dots, Z_n)$ . This means that the solution of  $\hat{T}_{b_n, \delta_n}$  by itself (with no threshold) is capable of detecting independence.

On the other hand, under  $\mathcal{H}_1$ , we notice that it is only relevant to bound the under-estimation of  $i_n(Z_1, \dots, Z_n)$  with the empirical (log-likelihood ratio) information  $\hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n))$ . Limiting the analysis only to the underestimation of  $i_n(Z_1, \dots, Z_n)$  can be

argued from the well-known observation that quantization (on average) provides a bias (under-estimation) estimation of  $\mathbb{E}(i_n(Z_1, \dots, Z_n)) = I(X; Y)$ . More importantly for our problem, an overestimation of the oracle information  $\hat{i}_n(Z_1, \dots, Z_n)$  does not increase the type 2 error when using  $\hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n))$  instead of the oracle  $\hat{i}_n(Z_1, \dots, Z_n)$  under  $\mathcal{H}_1$ .

In summary, Theorem 2 indicates that with an arbitrary high probability, the empirical TSP solution of (17) achieves the performance of the oracle solution that optimizes the upper bound on the regret derived in Eq.(13) over the TSP family  $\{\pi_T, T \ll T_b^{full}\}$ .

*Remark 1:* In the proof of Theorem 2, we determine a set  $\mathcal{E}_{\delta_n, b_n}^n$  (see Eq.(39)) where if  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^n$  then (20) holds under  $\mathcal{H}_0$  and (21) holds under  $\mathcal{H}_1$ . Importantly, we obtain that  $\mathbb{P}(\mathcal{E}_{\delta_n, b_n}^n) \geq 1 - \delta_n$  from Lemma 1. Therefore, in Theorem 2,  $N(\epsilon)$  is fully determined by  $(\delta_n)_{n \geq 1}$  for any  $\epsilon > 0$ . Indeed,  $N(\epsilon) = \inf\{n \geq 1, \text{st. } \delta_n < \epsilon\}$ , which is well-defined since  $(\delta_n)$  is  $o(1)$ . In particular, for the case  $(1/\delta_n)$  being  $\mathcal{O}(e^{n^{1/3}})$ , it follows that  $N(\epsilon)$  is  $\mathcal{O}(\ln(1/\epsilon)^3)$ .

### A. Detecting Independence With the Structure of $\hat{T}_{b_n, \delta_n}$

From (20), our data-driven partition has the capacity to detect independence (under  $\mathcal{H}_0$ ) in a stronger structural way: (with high probability) the data-driven partition collapses to the trivial cell, i.e.,  $\pi_{b_n, \delta_n} = \{\mathbb{R}^d\}$ , with a finite sample size. Here, we improve this result by showing that the condition  $\pi_{b_n, \delta_n} = \{\mathbb{R}^d\}$  happens within a finite sample size almost surely.

Let us introduce the (random) time at which the partition process  $(\pi_{b_n, \delta_n})_{n \geq 1}$  collapses to the trivial set  $A_{root} = \{\mathbb{R}^d\}$ :

*Definition 3:*  $\mathcal{T}_0((Z_n)_{n \geq 1}) \equiv$

$$\sup \left\{ m \geq 1 : \left| \hat{T}_{b_m, \delta_m} \right| > 1 \right\} \in \mathbb{N}^* \equiv \mathbb{N} \cup \{\infty\}. \quad (22)$$

If  $\mathcal{T}_0((Z_n)_{n \geq 1}) = k$ , it means that  $|\hat{T}_{b_n, \delta_n}| = 1 \Leftrightarrow \pi_{b_n, \delta_n} = \{\mathbb{R}^d\}$  for any  $n > k$ . Therefore, if  $\mathcal{T}_0((Z_n)_{n \geq 1}) \in \mathbb{N}$ , this value is expressing the last time the data-driven partition is different from the trivial solution  $\{\mathbb{R}^d\}$ . On the other hand, if  $\mathcal{T}_0((Z_n)_{n \geq 1}) = \infty$ , this condition means that non-trivial partitions are observed infinitely often (i.o.) in the sequence  $(\pi_{b_n, \delta_n}(Z^n))_{n \geq 1}$ . In general, we could have that  $\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) = \infty) > 0$  under  $\mathcal{H}_0$ , which does not contradict the capacity that  $(\phi_{b_n, \delta_n, a_n}(\cdot))_{n \geq 1}$  has for detecting independence consistently under  $\mathcal{H}_0$  (see Corollary 1).

The following result shows that  $\mathcal{T}_0((Z_n)_{n \geq 1})$  is finite with probability one under some mild conditions.

*Lemma 3:* Let us assume that  $(b_n) \approx (n^{-l})$  for  $l \in (0, 1/3)$  and  $(\delta_n)$  is  $\ell_1(\mathbb{N})$ ,  $(1/\delta_n)$  is  $\mathcal{O}(e^{n^{1/3}})$ . Then under the hypothesis that  $P \in \mathcal{P}_0$ :  $\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) < \infty) = 1$ .

The proof is presented in Appendix D.

Importantly, the condition  $\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) < \infty) = 1$  implies that  $(\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n))_{n \geq 1}$  reaches 0 eventually with probability one for any sequence  $(a_n)_{n \geq 1} \in [0, 1]^{\mathbb{N}}$ . From this observation, we obtain the following result that improves the regime of parameters where a consistent test is achieved (established in Corollary 1).

*Theorem 3:* If  $(b_n)_{n \geq 1} \approx (n^{-l})$  for  $l \in (0, 1/3)$ ,  $(\delta_n)_{n \geq 1}$  is  $\ell_1(\mathbb{N})$ ,  $(1/\delta_n)_{n \geq 1}$  is  $\mathcal{O}(e^{n^{1/3}})$ , and  $(a_n)_{n \geq 1}$  is  $o(1)$ , then the scheme  $(\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n))_{n \geq 1}$  is strongly consistent for detecting independence (Def. 1).

The proof is presented in Appendix E.

Finally, if we focus on the admissible solution where  $(1/\delta_n) \approx (e^{n^{1/3}})$ , we have the following refined description for the distribution of  $\mathcal{T}_0((Z_n)_{n \geq 1})$  under  $\mathcal{H}_0$ .

*Theorem 4:* Under the assumption of Theorem 3, let us consider that  $(1/\delta_n) \approx (e^{n^{1/3}})$ . Under  $\mathcal{H}_0$ , we have that  $\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \geq m) \leq K e^{-m^{1/3}} \sim \mathcal{O}(e^{-m^{1/3}})$  for any  $m \geq 1$  and for some universal constant  $K > 0$ .

The proof is presented in Appendix F.

### B. Remarks About Theorem 4

**1:** If we introduce  $\mathcal{T}((Z_n)_{n \geq 1}, (a_n)_{n \geq 1}) \equiv$

$$\sup \{ m \geq 1 : \phi_{b_m, \delta_m, a_m}(Z_1, \dots, Z_m) = 1 \}, \quad (23)$$

as the first time when the binary process  $(\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n))_{n \geq 1}$  reaches and stays at 0, it is simple to verify that  $\mathbb{P}(\mathcal{T}((Z_n)_{n \geq 1}, (a_n)_{n \geq 1}) \geq m) \leq \mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \geq m)$  for any  $(a_n)_n$  (see Appendix E). Therefore, under  $\mathcal{H}_0$  and the assumptions of Theorem 4, we have as a corollary that  $\mathbb{P}(\mathcal{T}((Z_n)_{n \geq 1}, (a_n)_{n \geq 1}) \geq m) \leq K e^{-m^{1/3}} \sim \mathcal{O}(e^{-m^{1/3}})$  for any sequence  $(a_n)_{n \geq 1}$  as long as  $(a_n)_{n \geq 1}$  is  $o(1)$ .

**2:** With this bound, we could determine (under  $\mathcal{H}_0$ ) a lower bound on the number of samples needed to guarantee that we detect independence structurally, i.e., reaching  $\pi_{b_n, \delta_n} = \{\mathbb{R}^d\}$ , and also with our schemes  $(\phi_{b_n, \delta_n, a_n}(\cdot))_{n \geq 1}$  with a confidence probability  $1 - \epsilon$ . This critical number of samples is achieved for  $m(\epsilon)$  such that  $K e^{-m(\epsilon)^{1/3}} < \epsilon$  and  $K e^{-(m(\epsilon)-1)^{1/3}} \geq \epsilon$ . This number scales like  $(m(\epsilon))_\epsilon \approx (\ln(K/\epsilon))^3 \sim \mathcal{O}(\ln(1/\epsilon)^3)$ .

## VII. IMPLEMENTATION

### A. The Learning Algorithm

Let us briefly revisit the stages presented in Section V for the construction of our TSP scheme. Beginning with the partition, there are two phases: growing and pruning. In the growing phase, presented in Section V-A, we have a collection of tree-indexed partitions  $\{\pi_T, T \ll T_b^{full}\}$  where  $\pi_T = \{A_v, v \in \mathcal{L}(T)\}$  and by construction  $\hat{P}_n(A_v) \geq b_n = w \cdot n^{-l}$  for any  $v \in T$ . Here, we added a scalar parameter  $w > 0$  as a relevant design element for our numerical analysis. In the pruning phase, we implement the following complexity-regularized information maximization:

$$\hat{T}_{b_n, \delta_n}(\alpha) \equiv \arg \max_{T \ll T_b^{full}} D_{\sigma(\pi_T)}(\hat{P}_n || \hat{Q}_n^*) - \alpha \cdot r_{b_n, \delta_n}(|T|), \quad (24)$$

where we included a regularized parameter  $\alpha \in \mathbb{R}^+$ .

## B. Computational Cost

To analyze the cost of implementing (24), we analyze the cost of each of its two phases individually.

**Growing Phase:** During the growing stage (see Section V-A), we obtain a collection of tree-indexed partitions  $\{\pi_T, T \ll T_{b_n}^{full}\}$ , where  $\pi_T = \{A_v, v \in \mathcal{L}(T)\}$ . The full tree  $T_{b_n}^{full}$  is obtained after a series of statistically equivalent binary partitions. The computational cost of the growing phase depends on the number of sub-partitions (or binary splittings) needed to obtain the full tree (which depends on the number of cells), and the cost of sorting the projected data of each cell and selecting the median of the ordered sequence to partition the cell. The number of cells in the full tree  $|\pi_{T_{b_n}^{full}}|$  can be estimated considering that our design criterion states that  $\hat{P}_n(A_v) \geq b_n = w \cdot n^{-l}$  for any  $v \in T$ , where  $w \in \mathbb{R}^+$  and  $l \in (0, 1/3)$  are design parameters of our method. Then, we should have at least  $b_n \cdot n$  samples per cell. Therefore, the number of cells associated with  $T_{b_n}^{full}$  is bounded by

$$\left| \pi_{T_{b_n}^{full}} \right| \leq \frac{n}{b_n \cdot n} = \frac{1}{b_n} \approx n^l \quad (25)$$

On the other hand, the cost of sorting the projected data of each cell depends on the number of data points to be ordered. In particular, if a cell has  $m$  points, then the sorting cost is  $\mathcal{O}(m \log(m))$  [49]. As we reach deeper levels in the tree, the number of points per cell reduces. If we consider a level of depth  $k$  (where  $k = 1$  corresponds to the root and  $k = \lceil \log(|\pi_{T_{b_n}^{full}}|) \rceil \approx \lceil \log(n^l) \rceil$  corresponds to the full tree), there are  $2^{k-1}$  cells and each cell contains roughly  $n/2^k$  points. Therefore, the order of the total number of operations during the growing phase can be approximated to

$$\mathcal{O} \left( \sum_{k=1}^{\lceil \log(n^l) \rceil} 2^{k-1} \cdot \frac{n}{2^{k-1}} \log \left( \frac{n}{2^{k-1}} \right) \right) = \mathcal{O}(n \log^2(n))$$

**Pruning Phase:** The main regularized info-max problem in Eq. (24) is solved in this stage. The first term  $D_{\sigma(\pi_T)}(\hat{P}_n || \hat{Q}_n^*)$  is additive (in the sense presented in [45]), whereas the penalty  $r_{b_n, \delta_n}(|T|)$  scales like  $\mathcal{O}(\sqrt{|T|})$ , and, therefore, it is a sub-additive function of the size of the tree. In this scenario, Scott [45, Th. 1] showed that the family of minimum cost trees  $\{R_i\}_{i=1}^m$ :

$$R_i \equiv \arg \max_{\substack{T \ll T_{b_n}^{full} \\ |T|=m-i+1}} D_{\sigma(\pi_T)}(\hat{P}_n || \hat{Q}_n^*) \quad , i \in \{1 \dots m\} \quad (26)$$

is embedded, i.e.,  $R_1 = T_{b_n}^{full} \gg R_2 \gg \dots \gg R_m = \{A_{root}\}$ . Furthermore, [45, Th. 2] states that the solution of our problem in (24) corresponds to one of these embedded trees: i.e., for any  $\alpha > 0$ ,  $\exists i \in \{1 \dots m\}$  such that  $\hat{T}_{b_n, \delta_n}(\alpha) = R_i$ . This embedded property allows computationally efficient algorithms to be designed. Indeed, Scott [45] presented two algorithms to solve (24) with a worst case cost of  $\mathcal{O}(|T_{b_n}^{full}|^2)$ . In our case we have that  $|T_{b_n}^{full}| \leq 1/b_n$ , then the pruning stage in (24) has a computational cost of  $\mathcal{O}(n^{2/3})$ , which is a polynomial (sub-linear) function of  $n$ . The pseudo code of this stage is presented in Algorithm 1.

---

### Algorithm 1: TSP Family Pruning .

---

**Input :** full tree  $T_{b_n}^{full}$ , minimum probability of partitions  $b_n$ , confidence probability  $\delta_n$ , regularization parameter  $\alpha$

**Output:** pruned tree  $\hat{T}_{b_n, \delta_n}$  in Eq.(24)

Initialization;

$\hat{T}_{b_n}^1 = \{root\}$ ;

Family pruning problem of embedded trees;

**for**  $k = 2$  **to**  $|T_{b_n}^{full}|$  **do**

Maximum empirical information tree of size  $k$ ;

$v^* = \arg \max_{v \in \mathcal{L}(\hat{T}_{b_n}^{k-1})} D_{\sigma}(\pi_{\hat{T}_{b_n}^{k-1} \cup \{l(v), r(v)\}})(\hat{P}_n || \hat{Q}_n^*)$ ;

$\hat{T}_{b_n}^k = \hat{T}_{b_n}^{k-1} \cup \{l(v^*), r(v^*)\}$ ;

**end**

Regularized empirical information maximization;

$\hat{T}_{b_n, \delta_n} =$

$\arg \max_{T \in \left\{ \hat{T}_{b_n}^1, \dots, \hat{T}_{b_n}^{|T_{b_n}^{full}|} \right\}} D_{\sigma(\pi_T)}(\hat{P}_n || \hat{Q}_n^*) - \alpha \cdot r_{b_n, \delta_n}(|T|)$ ;

$// r_{b_n, \delta_n}(|\hat{T}_{b_n}^1|) = 0$

---

In summary, the construction of our partition (growing and pruning) has a computational cost that is  $\mathcal{O}(n \log^2(n))$  on the sample size  $n$ .

## VIII. EMPIRICAL ANALYSIS

In this section, we present some controlled synthetic scenarios to evaluate the performance of our method. We begin analyzing how the selection of parameters affects the capacity of our scheme to estimate MI. From these results and insights, we analyze the performance of the solutions  $\{\hat{T}_{b_n, \delta_n}(\alpha), \alpha \geq 0\}$  for testing independence. For the rest of this section,  $\pi_{b_n, \delta_n}^\alpha(\cdot)$  denotes the TS partition,  $\hat{i}_{b_n, \delta_n}^\alpha(\cdot)$  denotes the MI estimator, and  $\phi_{b_n, \delta_n, a_n}^\alpha(\cdot)$  denotes the final test.

### A. Testing Independence: Non-Asymptotic Empirical Analysis

We analyze the problem of testing independence with our data-driven test  $\phi_{b_n, \delta_n, a_n}^\alpha(\cdot)$  in (7). Here we focus on the non-asymptotic capacity of our framework to detect the two hypotheses. For this, we propose the following detection times:<sup>10</sup>

$$\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1}) \equiv \sup \{m \geq 1 : \phi_{b_m, \delta_m, a_m}^\alpha(Z_1, \dots, Z_m) = 1\}, \quad (27)$$

$$\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1}) \equiv \sup \{m \geq 1 : \phi_{b_m, \delta_m, a_m}^\alpha(Z_1, \dots, Z_m) = 0\}. \quad (28)$$

$\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$  are random variables (rvs.) in  $\mathbb{N}^* = \mathbb{N} \cup \{\infty\}$  determining when our test reaches 0 and 1, respectively. Indeed, Theorem 3 (under specific conditions) tells us that under  $\mathcal{H}_0$ ,  $\mathbb{P}(\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1}) < \infty) = 1$  and that under  $\mathcal{H}_1$ ,

<sup>10</sup>To simplify the notation, the dependency of  $\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$  on  $(b_m)$ ,  $(\delta_m)$ ,  $(a_m)$  and the scalar parameters  $w$ ,  $\alpha$  introduced in Section VII-A will be considered implicit.



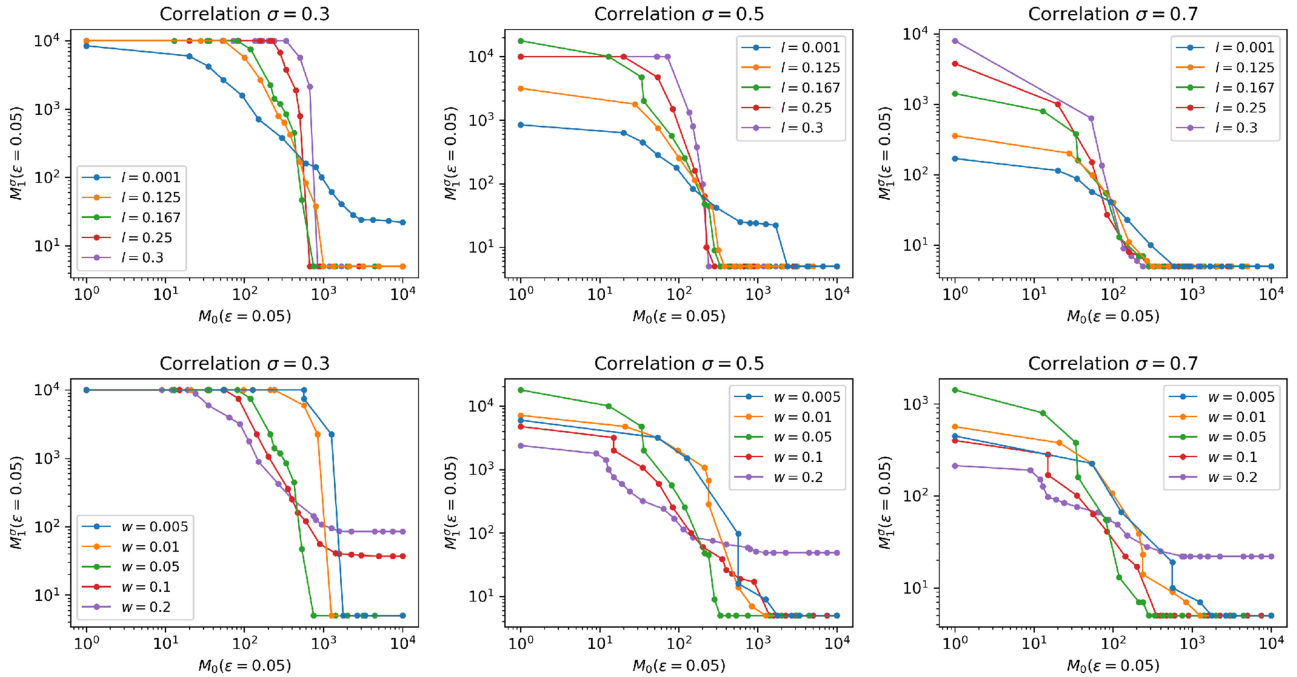


Fig. 1. Illustration of the trade-off between  $M_0(\epsilon = 0.05)$  and  $M_1^\sigma(\epsilon = 0.05)$ , presented in Eqs. (29) and (30), for different parameters of our TSP approach. Each curve is associated with fixed values of  $l$  and  $w$  and is produced using different values of  $\alpha$ . Three scenarios of correlation (under  $\mathcal{H}_1$ ) are  $\sigma = \mathbb{E}(XY) \in \{0.3, 0.5, 0.7\}$ . For the graph in the first row  $w = 0.05$ , whereas  $l = 0.167$  in the graph depicted in the second row.

$\mathbb{P}(\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1}) < \infty) = 1$ . However, we are interested in the complete distribution of the rvs.  $\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$  under  $\mathcal{H}_0$  and  $\mathcal{H}_1$ , respectively. In particular, we are interested in evaluating the pmf of  $\tilde{\mathcal{T}}_i(\cdot)$ , i.e.,  $(\mathbb{P}(\tilde{\mathcal{T}}_i((Z_n)_{n \geq 1}) = k))_{k \geq 1}$  under  $\mathcal{H}_i$  (with  $i \in \{0, 1\}$ ). Looking at these distributions, for any  $\epsilon > 0$  we can define

$$M_0(\epsilon) \equiv \min \left\{ m \geq 1, \mathbb{P}(\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1}) \leq m) \geq 1 - \epsilon \right\} \quad (29)$$

$$M_1(\epsilon) \equiv \min \left\{ m \geq 1, \mathbb{P}(\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1}) \leq m) \geq 1 - \epsilon \right\} \quad (30)$$

$M_0(\epsilon)$  (and  $M_1(\epsilon)$ ) indicates how many observations (sampling complexity) are needed to detect independence (and statistical dependency) with a confidence probability  $1 - \epsilon$  under  $\mathcal{H}_0$  (and  $\mathcal{H}_1$ ). In the context of our solution, for a given TSP scheme (function of the parameters  $l$ ,  $w$ ,  $\alpha$  and the sequence  $(a_n)$ ), we will look at the trade-off expressed in the pair  $(M_0(\epsilon), M_1(\epsilon))$  of induced tests when varying key parameters of our method.

1) *Simulation Setting*: We consider a joint vector  $Z = (X, Y)$  following a zero-mean Gaussian distribution in  $\mathbb{R}^2$  where the correlation coefficient determining  $I(X, Y)$  is parametrized by  $\sigma = \mathbb{E}(XY)$ . Concerning the alternative hypothesis ( $\mathcal{H}_1$ ), we consider different levels of MI indexed by  $\sigma = \mathbb{E}(XY) \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ . We use 1,000 iid realizations of  $(X, Y)$  under the different scenarios ( $\mathcal{H}_0$  and  $\mathcal{H}_1^\sigma$ ) to run our test with these iid samples. By doing so, we obtain 1,000 realizations of the rvs.  $\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$  and with those realizations we obtain an empirical estimation of their pmfs. and

empirical estimations of  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  (indexed by the value  $\sigma$ ), respectively.<sup>11</sup>

2) *Parameter Selection for  $\phi_{b_n, \delta_n, a_n}^\alpha(\cdot)$* : To evaluate the sensitivity of our TSP scheme, we consider the following fixed sequences  $(a_n) = (0.5 \cdot n^{-1})$  and  $(\delta_n) = (e^{-n^{1/3}})$  in the admissible regime established in Theorem 3.<sup>12</sup> Considering  $(b_n) = (w \cdot n^{-l})$  (parametrized by  $w$  and  $l$ ) and  $\alpha$  (used to solve  $\hat{T}_{b_n, \delta_n}(\alpha)$  in Eq.(24)), we consider a preliminary analysis on MI estimation to select a range of reasonable values (in Supplemental Material–Section I). In particular, we consider  $l \in \{0.001, 0.125, 0.167, 0.25, 0.3\}$  and  $w \in \{0.005, 0.01, 0.05, 0.1, 0.2\}$ . We proceed as follows: given fixed parameters  $l$  and  $w$ , we explore values of  $\alpha$  in  $[0, 4 \cdot 10^{-4}]$  to express the trade-off between  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  under different data scenarios:  $\sigma \in \{0.3, 0.5, 0.7\}$ .

3) *Results*: Fig. 1 shows the curves expressing the trade-off between  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  for different parameter configurations of our method in the range for  $l$  and  $w$  mentioned above and for  $\epsilon = 0.05$ . In general, we notice that the effect of these two parameters is relevant in the trade-off expressed in the curves. Both  $w$  and  $l$  determine the growing phase (i.e., the creation of the full tree) and also the pruning phase (regularization in Eq.(24)) because  $r_{b_n, \delta_n}(|T|)$  is also a function of  $b_n$ . Therefore,

<sup>11</sup>By the law of large numbers, the estimators of  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  are strongly consistent. 1000 samples shown to be sufficient for the purpose of the analysis.

<sup>12</sup> $(a_n)_{n \geq 1}$  needs to be  $o(1)$  and  $(\delta_n)$  needs to be  $\ell_1(\mathbb{N})$ . Other configurations for  $(a_n)$  and  $(\delta_n)$  can be explored within the admissible range declared in Theorem 3.




Fig. 2. Empirical estimation of the probability mass functions of  $\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  in Eq.(27) and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$  in Eq.(28). Four scenarios of the regularization parameter  $\alpha$  are considered to illustrate its effects in the distribution of  $\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$ . For each case, the values of  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  are also illustrated (vertical lines) considering  $\epsilon = 0.05$  and  $\sigma = 0.3$ .

the effects of  $(b_n)_{n \geq 1}$  on the results are not easy to express theoretically.

Describing the curves, we could say in general that a bigger value of  $w$  within the explored range reduces the full tree's size in the growing phase. This effect is expressed in a better trade-off of the pair  $(M_0(\epsilon), M_1^\sigma(\epsilon))$  in the regime where  $M_0(\epsilon) \leq 10^2$ , at the expense of a worse trade-off of  $(M_0(\epsilon), M_1^\sigma(\epsilon))$  when  $M_0(\epsilon) \geq 10^3$ . A similar general effect in the trade-off  $(M_0(\epsilon), M_1^\sigma(\epsilon))$  is produced when decreasing the value of  $l$  within the explored range. Our family of solutions offers a collection of different trade-offs between  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  by exploring different values of  $\alpha \in [0, 4 \cdot 10^{-4})$  in our solution. Therefore, the selection of the best parameters should be a function of the regime of sample size we want to consider for the detection of  $\mathcal{H}_0$ . For the final comparison with alternative approaches, we decided to consider one of the curves with a less prominent decreasing transition in Fig. 1 (obtained for  $w = 0.1$  and  $l = 0.001$ ).

Fig. 2 illustrates the estimated (empirical) pmfs of  $\tilde{\mathcal{T}}_0((Z_n)_{n \geq 1})$  and  $\tilde{\mathcal{T}}_1((Z_n)_{n \geq 1})$  under  $\sigma = 0$  ( $\mathcal{H}_0$ ) and  $\sigma = 0.3$  ( $\mathcal{H}_1$ ), respectively. These histograms illustrate how the distributions are affected by the regularization parameter  $\alpha$  and, with that, the trade-off in  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  observed in Fig. 1. Specifically, a small value of  $\alpha$  concentrates the distribution of  $\mathcal{T}_0((Z_n)_{n \geq 1})$  in a high number of observations, while  $\mathcal{T}_1((Z_n)_{n \geq 1})$  concentrates on a lower number of observations. This leads to a high value of  $M_0(\epsilon)$  and a low value of  $M_1^\sigma(\epsilon)$  for that specific  $\alpha$ . The opposite holds for higher values of  $\alpha$ .

Finally, we compared our TSP scheme in terms of the trade-off in  $(M_0(\epsilon), M_1^\sigma(\epsilon))$  with  $\epsilon = 0.05$  and  $\sigma \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$  (for  $\mathcal{H}_1$ ) exploring three strategies presented in the literature [20]: the  $L_1$  test, the log-likelihood test, and the Pearson- $\chi^2$  test. Each trade-off curve is generated by multiplying a well-selected range of values  $C \in \mathbb{R}$  to the corresponding threshold for the independence detection [20]. For the parameters of the alternative tests, we consider regimes where these tests are strongly consistent<sup>13</sup>. Within that, we selected

<sup>13</sup>We consider the number of partitions per dimension as  $m(n) = n^p$  with  $p \in (0, 0.5)$ , which satisfies the strong-consistency conditions of the  $L_1$  and the log-likelihood tests established in [20]. Although there is no explicit proof of Pearson- $\chi^2$  test strong-consistency, we use a regime of values for  $p \in (0, 1/3)$  according to the range suggested in [20].

parameters that offer a smooth transition between their sampling complexity pairs' trade-off as we did to select the parameters of our method.

Fig. 3 presents these curves for our method and the alternative approaches. These curves express the trade-off between the ability to detect independence and non-product probabilistic structure ( $\mathcal{H}_0$  and  $\mathcal{H}_1$ ) from data. In general, we could say that our method shows better results in almost all explored scenarios ( $\sigma \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ ) where the TSP trade-off curves dominate the others. There is one exception to this general trend in the regime of  $M_0(\epsilon) < 10^2$  for lower values of correlation for  $\mathcal{H}_1$ . However, in all the other regimes, our method performs better than its closest relative, which is the log-likelihood test that uses non-adaptive partitions [20]. Interestingly, our method's performance improvements increase with the magnitude of  $\sigma$  (for  $\mathcal{H}_1$ ). This shows that our approach's advantage is more prominent when the alternative scenario has a higher level of mutual information.

## B. Multidimensional Analysis

We conclude this analysis by evaluating our method in higher dimensions. As in the previous analysis  $d = 2$ , we need to select the parameters  $w$  and  $l$  of our scheme. In this multi-dimensional context, we select a value of  $l$  that offers a smooth trade-off curve according to the experimental results in Fig. 1. However, as  $w$  has a stronger impact on the resulting number of partitions, rather than choosing it individually for every dimension  $d$ , we propose  $w = C^{p \cdot d}$  as a heuristic rule for selecting  $w$  according to  $d$ , where  $C > 0$  and  $p > 0$ .

This rule and its parameters were designed on the principle that every dimension of the joint space  $\mathbb{R}^d$  should be explored at least once in the growing phase. This criterion ensures that our data-driven scheme explores every coordinate of  $\mathbb{R}^d$  in the pursuit of detecting relevant statistical dependencies under  $\mathcal{H}_1$ . For a dimension  $d > 0$ , the full tree needs to have at least  $2^d$  leaves (cells). For this, a basic condition is that  $b_n$  is lower bounded by  $1/n$  (our stopping criterion requires at least one sample per cell), which implies that to explore all the  $d$  dimensions we need at least  $2^d$  i.i.d. samples to meet this requirement. Then assuming the non-trivial regime  $n \geq 2^d$ , we selected  $C$  and  $p$  in

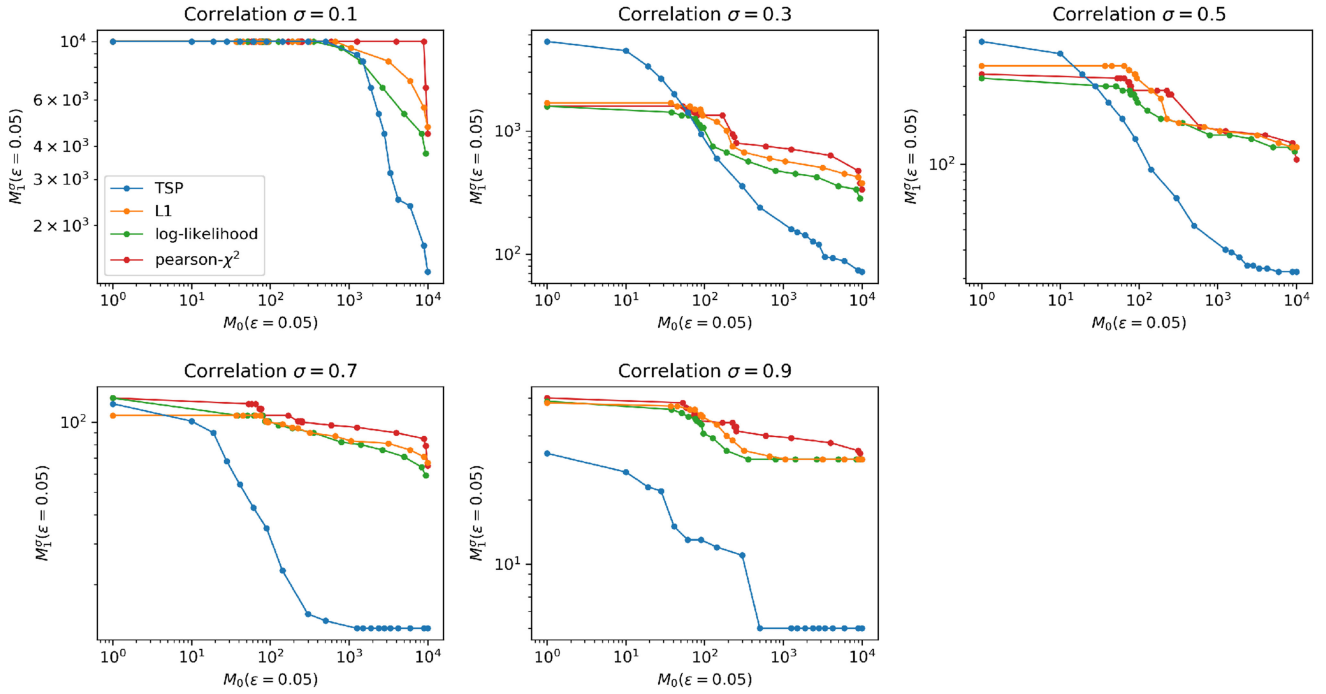


Fig. 3. Illustration of the trade-off between  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  obtained for our TSP test, the  $L_1$ -test, the log-likelihood test, and the  $\chi^2$ -test. The TSP test parameter configuration is  $l = 0.001$  and  $w = 0.1$ , with  $\alpha \in [0.00001, 0.0005]$ . The  $L_1$  test parameter configuration is  $m(n) = n^{0.2}$ , with  $C \in [0.7, 1.46]$ . The log-likelihood test parameter configuration is  $m(n) = n^{0.2}$ , with  $C \in [0.05, 0.3]$ . The Pearson- $\chi^2$  test parameter configuration is  $m(n) = n^{0.25}$ , with  $C \in [0.005, 0.185]$ . Five scenarios of correlation are presented for  $\mathcal{H}_1$  considering  $\sigma = \mathbb{E}(XY) \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ .




Fig. 4. Number of partitions generated for the full tree of the TSP under  $l = 0.001$  as a function of the dimensionality  $d$  of the joint space of the random variables for  $n \in \{10^3, 10^4, 10^5\}$  number of samples and two heuristic rules for the parameter  $w$ . Left:  $w = 0.1^{d/2}$ . Right:  $w = 0.225^{d/2}$ . Diagonal line determines the minimum number of partitions that allow each dimension of the joint space to be partitioned at least once.

our rule with the objective that  $|T_{b_n}^{full}| \geq 2^d$  is met independent of the value of  $d$  (as long as  $n \geq 2^d$ ). This happens if  $w$  becomes smaller when  $d$  increases, which happens in our rule for  $C < 1$  and  $p > 0$ .

Fig. 4 shows the number of cells of the full-tree, namely  $|T_{b_n}^{full}|$ , for two settings:  $w = 0.1^{d/2}$  ( $C = 0.1$  and  $p = 1/2$ ) in the left panel, and  $w = 0.225^{d/2}$  ( $C = 0.225$  and  $p = 1/2$ ) in the right panel. The curves are shown across  $d$  for three scenarios of sample-size  $n = 10^3$ ,  $n = 10^4$  and  $n = 10^5$ . We also include the curve  $2^d$  (dashed line) to indicate the critical lower bound. Both settings satisfy the proposed criterion of partitioning each dimension at least once (points above the diagonal) in the whole regime when  $n \geq 2^d$ . The horizontal line shows the points where the number of partitions generated by our method is equal

to the number of samples (which happens when  $b_n = 1/n$ ). Interestingly, when  $d$  approaches the critical condition  $2^d = n$ , our full tree meets a scenario where every cell has one sample. Of the two settings, the selection on the right panel is more conservative, and it is the one we will use for the rest of this analysis. This selection partitions each dimension just above the minimum requirement (for the values of  $d$  where  $n \geq 2^d$ ) and achieves the critical condition  $n = 2^d$  as  $d$  increases.

Using this rule  $w = 0.225^{d/2}$  and  $l = 0.001$ , we consider a zero-mean joint Gaussian vector  $(X, Y)$  in  $\mathbb{R}^d$ , where  $X = (X_1, \dots, X_p)$  and  $Y = (Y_1, \dots, Y_q)$  and, under  $\mathcal{H}_1$ , a diagonal covariance of the form  $cov(X_i, Y_j) = \delta_{ij} \sigma(d)^2$  (with  $\delta_{ij}$  the Kronecker's delta). For these results, we consider  $n = 10^4$  i.i.d samples meeting the condition that  $n \geq 2^d$  for  $d \in \{4, 6, 8, 10, 12\}$ .




Fig. 5. Illustration of the trade-off between  $M_0(\epsilon)$  and  $M_1^\sigma(\epsilon)$  obtained for the TSP test using samples from a multidimensional Gaussian distribution with component-wise correlation  $\text{cov}(X_i, Y_j) = \delta_{ij}\sigma(d)$ , where  $\sigma(d)$  is selected to preserve the mutual information between the variables along the dimensions  $p = q \in \{2, 3, 4, 5, 6\}$  of each random variable under  $l = 0.001$  and  $w = 0.225^{d/2}$  with  $d = p + q$ . Three scenarios of mutual information are explored  $I(X; Y) \in \{0.06803, 0.20752, 0.48572\}$ , which correspond to the theoretical mutual information obtained in the univariate Gaussian case  $p = q = 1$  with  $\mathbf{E}(XY) = \sigma, \sigma \in \{0.3, 0.5, 0.7\}$ , respectively.

For the alternative, we consider  $\sigma \in \{0.3, 0.5, 0.7\}$  for  $d = 2$  ( $p = q = 1$ ), and we choose the values of  $\sigma(d)$  for  $d \geq 4$  such that the mutual information between  $X$  and  $Y$  is the same as that obtained for  $d = 2$  with  $\sigma \in \{0.3, 0.5, 0.7\}$ . Therefore, for the alternative, we create three scenarios of constant MI  $I(X; Y) \in \{0.06803, 0.20752, 0.48572\}$  across dimensions. This experimental design allows us to isolate the effect of dimensionality from the discrimination of the underlying task, which is known to be proportional to the value of MI under  $\mathcal{H}_1$  [32], [35], [36].

Fig. 5 presents the trade-off between  $(M_0(\epsilon), M_1^\sigma(\epsilon))$  for the following values of  $\sigma(d)$  (under  $\mathcal{H}_1$ ) and for different dimensions  $d \in \{4, 6, 8, 10, 12\}$ . First, we confirm that our scheme detects both hypotheses with a finite sample size as our theoretical results predict. Here, we observe that by increasing the dimensions of the problem the curves  $(M_0(\epsilon), M_1^\sigma(\epsilon))$  show sharper transitions (beyond a critical value for  $M_0(\epsilon)$ ) and overall the problems become (in general) more difficult: for a given value of  $M_0(\epsilon)$ , its respective  $M_1^\sigma(\epsilon)$  increases with the dimension  $d = p + q$ . In other words, for a relatively good capacity to detect independence (when  $M_0(\epsilon) \leq 10^3$ ), it is more challenging for our test to detect the alternative as  $d$  is higher. This performance trend could be attributed to the observation that in higher dimensions it is more challenging for a non-parametric framework to detect salient features under the alternative  $\mathcal{H}_1$  and, consequently, more observations (evidence) are required to correctly detect the alternative with probability  $1 - \epsilon$ . On the other hand, if we fix a dimension, let us say  $d = p + q = 8$ , the higher the MI the better the trade-off between  $(M_0(\epsilon), M_1^\sigma(\epsilon))$ , which is a trend consistent with our observation that MI is an indicator of the task discrimination, and consequently, the performance of our scheme is sensitive to the level of MI under the alternative.

## IX. SUMMARY AND FINAL DISCUSSION

We present a novel framework to address the problem of universal testing of independence using data-driven representations. Consistency and finite length (non-asymptotic) results express our solution's capacity to adapt its representations to the problem's sufficient statistics. This capacity is particularly

evident under the hypothesis of independence. Precise results show that our scheme detects this scenario – collapsing to the trivial partition with one cell for representing the joint space – with a finite sample size. On the experimental side, our solution offers a computationally efficient implementation.

In a controlled experimental setting, we show that our scheme offers a clear advantage compared with alternative non-parametric solutions that do not adapt their representations to the data. These results confirm the critical role that data-driven partitions could play in implementing a universal test of independence. Further analysis is needed to fully uncover this method's potential for machine learning (ML) problems and other applications. We anticipate that ML algorithms could benefit from both the representations obtained from our solutions (as an approximator of sufficient statistics) and our solution's capacity to detect independence with a finite number of samples.

### A. Limitations and Future Work

Our distribution-free results in Section VI provide a range of admissible parameters for our test; however, they do not provide a criterion for selecting them in a specific problem setting. To address this last practical aspect, we select a set of parameters from empirical observations and some basic heuristic rules. Then, it is an open research avenue to fully study conditions that would make a selection of optimal parameters (or nearly optimal) for our test under some specific conditions. Along these lines, it is relevant to further investigate specific classes of models, and based on this model assumption find a more constrained range of good parameters with improved finite-length performance results. In this vein, there are interesting ideas and results worth exploring about adaptivity that have been explored in statistical learning [50] and universal source coding [51], [52].

### B. A Recent Related Work

The authors in [34] have recently introduced a non-parametric test with a similar oracle-base design principle. They proposed a data-adaptive weight strategy to approximate the ideal weights of a NP statistic. That strategy has an interesting connection with our design approach in Section IV because they adapt

key parameters of their method to approximate a best (oracle) statistic. On the differences, the authors in [34] addressed the independence test by performing a binary expansion filtration of the observations and adapting the weights matrix of a quadratic form. Instead, our work addresses the independence test by estimating the mutual information (our statistics) through an adaptive tree-structured partition of the observation space. Therefore, although both approaches share a similar learning principle, the methods and components used to adapt them to the sample are entirely different.

### C. Statistical Significance Analysis

We want to conclude with a discussion about the statistical significance of our test. Given our test  $\phi_{b_n, \delta_n, a_n}(\cdot)$ , and assuming  $\mathcal{H}_0$ , the statistical significance is expressed by  $\alpha(b_n, \delta_n, a_n) \equiv \mathbb{P}(\{\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n) = 1\})$ . Importantly, from the proofs of Theorem 2 and Theorem 3, we have the following:

*Proposition 2:* Under the assumptions of Theorem 2 on  $(b_n)$  and  $(\delta_n)$ , for any  $P \in \mathcal{P}_0$  ( $\mathcal{H}_0$ ) and any  $a_n > 0$ , it follows that  $\alpha(b_n, \delta_n, a_n) \leq \delta_n, \forall n \geq 1$ .<sup>14</sup>

This result implies that the statistical significance of our test is fully controlled by the sequence  $(\delta_n)$ , which is one of our design parameters. Furthermore, the proof of this result shows that this bound is uniform over the class of models in  $\mathcal{P}_0$ . Direct corollaries of Proposition 2 are the following:

- adding the assumptions of Theorem 3 and under  $\mathcal{H}_0$ , it follows that  $\lim_{n \rightarrow \infty} \alpha(b_n, \delta_n, a_n) = 0$ . Importantly, this convergence to zero is uniform for every model in  $\mathcal{P}_0$ . [Uniform vanishing condition on  $(\alpha(b_n, \delta_n, a_n))_{n \geq 1}$ ]
- adding the assumptions of Theorem 4 and under  $\mathcal{H}_0$ , it follows that for any  $n \geq 1$ ,  $\alpha(b_n, \delta_n, a_n) \leq \delta_n \approx e^{-n^{1/3}}$ . Therefore, we achieve an exponentially fast vanishing error under  $\mathcal{H}_0$ . Importantly, this velocity is obtained uniformly (distribution-free) over  $P \in \mathcal{P}_0$ .

These uniform bounds on the statistical significance are obtained under  $\mathcal{H}_0$ . Under  $\mathcal{H}_1$ , we have from strong consistency (Definition 1) that  $\mathbb{P}(\{\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n) = 1\})$  (the power of the test) tends to 1 as  $n$  tends to infinity.<sup>15</sup>

#### APPENDIX A PROOF OF THEOREM 1

*Proof:* In general, first we know that if  $(\phi_n(Z_1^n))_{n \geq 1}$  reaches 0 eventually with probability one, it is equivalent to saying that the process  $(\phi_n(Z_1^n))_{n \geq 1}$  does not visit the event  $\{1\}$  infinitely often (i.o.). This observation reduces to verify the following

$$\begin{aligned} & \mathbb{P} \left( \limsup_n \{z_1^n \in \mathbb{R}^{dn} : \phi_n(z_1^n) = 1\} \right) \equiv \\ & \mathbb{P} \left( \bigcap_{m \geq 1} \bigcup_{n \geq m} \{z_1^n \in \mathbb{R}^{dn} : \phi_n(z_1^n) = 1\} \right) = 0. \end{aligned} \quad (31)$$

Equivalently, saying that  $(\phi_n(Z_1^n))_{n \geq 1}$  reaches 1 eventually with probability one reduces to

$$\begin{aligned} & \mathbb{P} \left( \limsup_n \{z_1^n \in \mathbb{R}^{dn} : \phi_n(z_1^n) = 0\} \right) \equiv \\ & \mathbb{P} \left( \bigcap_{m \geq 1} \bigcup_{n \geq m} \{z_1^n \in \mathbb{R}^{dn} : \phi_n(z_1^n) = 0\} \right) = 0. \end{aligned} \quad (32)$$

For the rest of the proof, we use  $O_1^n, \dots, O_n^n$  to denote  $O_{\pi_n}(Z_1), \dots, O_{\pi_n}(Z_n)$  and  $o_1^n, \dots, o_n^n$  to denote  $O_{\pi_n}(z_1), \dots, O_{\pi_n}(z_n)$ .

Under  $\mathcal{H}_0$ , we have from the third hypothesis that  $\hat{i}_{\pi_n}(O_1^n, \dots, O_n^n)/a_n$  tends to 0 with probability one, which means that for any  $\epsilon > 0$

$$\mathbb{P} \left( \bigcap_{m \geq 1} \bigcup_{n \geq m} B_{\epsilon, n}^{\pi_n, a_n} \right) = 0, \quad (33)$$

where  $B_{\epsilon, n}^{\pi_n, a_n} \equiv \{z_1^n \in \mathbb{R}^{dn} : \hat{i}_{\pi_n}(o_1^n, \dots, o_n^n) > \epsilon \cdot a_n\}$ . Under  $\mathcal{H}_0$ , let us consider the error event of  $\phi_{\pi_n, a_n}(\cdot)$  by

$$E_n^{\pi_n, a_n} \equiv \{z_1^n \in \mathbb{R}^{dn} : \phi_{\pi_n, a_n}(o_1^n, \dots, o_n^n) = 1\}.$$

If  $z^n \in E_n^{\pi_n, a_n}$  (by definition of the rule  $\phi_{\pi_n, a_n}(\cdot)$ ), it follows that  $\hat{i}_{\pi_n}(o_1^n, \dots, o_n^n) \geq a_n$  (see (7)). Then from definition of  $B_{\epsilon, n}^{\pi_n, a_n}$ , for any  $\epsilon < 1$  it follows that  $E_n^{\pi_n, a_n} \subset B_{\epsilon, n}^{\pi_n, a_n}$ . This implies that  $\bigcap_{m \geq 1} \bigcup_{n \geq m} E_n^{\pi_n, a_n} \subset \bigcap_{m \geq 1} \bigcup_{n \geq m} B_{\epsilon, n}^{\pi_n, a_n}$  which from (33) implies that

$$\mathbb{P} \left( \bigcap_{m \geq 1} \bigcup_{n \geq m} E_n^{\pi_n, a_n} \right) = 0. \quad (34)$$

On the other hand, under the assumption that  $I(X, Y) > 0$  ( $\mathcal{H}_1$ ), let us choose  $\epsilon_0 \in (0, I(X, Y))$ . Then, from the second assumption (see Definition 2), it follows that  $\hat{i}_{\pi_n}(O_1^n, \dots, O_n^n)$  converges to something strictly greater than  $\epsilon_0$  with probability one. This formally means that

$$\mathbb{P} \left( \bigcap_{m \geq 1} \bigcup_{n \geq m} \{z_1^n \in \mathbb{R}^{dn} : \hat{i}_{\pi_n}(o_1^n, \dots, o_n^n) \in (0, \epsilon_0)\} \right) = 0, \quad (35)$$

The error event of  $\phi_{\pi_n, a_n}(\cdot)$  under  $\mathcal{H}_1$  in this case is

$$\tilde{E}_n^{\pi_n, a_n} \equiv \{z_1^n \in \mathbb{R}^{dn} : \hat{i}_{\pi_n}(o_1^n, \dots, o_n^n) < a_n\}. \quad (36)$$

Finally using the condition that  $(a_n)_{n \geq 1}$  tends to zero with  $n$  ( $(a_n)_n$  is  $o(1)$ ), we have that

$$\tilde{E}_n^{\pi_n, a_n} \subset \{z_1^n \in \mathbb{R}^{dn} : \hat{i}_{\pi_n}(o_1^n, \dots, o_n^n) \in (0, \epsilon_0)\} \quad (37)$$

eventually in  $n$ , which implies from (35) that

$$\mathbb{P} \left( \bigcap_{m \geq 1} \bigcup_{n \geq m} \tilde{E}_n^{\pi_n, a_n} \right) = 0. \quad (38)$$

□

#### APPENDIX B PROOF OF LEMMA 2

We use the following results from [39].

<sup>14</sup>The proof is presented in the Supplementary Material.

<sup>15</sup>This argument is presented in the Supplementary Material.

*Lemma 4:* [39, Th.3] Under the assumptions of Lemma 2,  $\lim_{n \rightarrow \infty} \hat{i}_{\pi_{b_n, \delta_n}}(O_1, \dots, O_n) = I(X, Y)$  with probability one.

*Lemma 5:* [39, Th.4] Under the assumptions of Lemma 2 part ii), for any  $p > 0$ ,  $\hat{i}_{\pi_{b_n, \delta_n}}(O_1, \dots, O_n)$  is  $o(n^{-p})$ .<sup>16</sup>

*Proof:* Eq.(18) follows from Lemma 4, the fact that by the strong law of large number  $\lim_{n \rightarrow \infty} i_n(Z_1, \dots, Z_n) = I(X, Y)$  with probability one and the union bound.

Under the assumption that  $P \in \mathcal{P}_0$  (i.e.  $I(X, Y) = 0$ ), we have that  $P = Q^*(P)$ . This implies by definition that  $i_n(Z_1, \dots, Z_n) = 0$  with probability one. Then the regret is basically our empirical information term  $\hat{i}_{\pi_{b_n, \delta_n}}(O_1, \dots, O_n)$ . Then, Lemma 5 implies the result in (19).  $\square$

#### APPENDIX C PROOF OF THEOREM 2

*Proof:* If we define the event  $\mathcal{E}_{\delta_n, b_n}^{n, k} \equiv$

$$\left\{ z_1^n : \sup_{T \in \mathcal{G}_{b_n}^k} \left| D_{\sigma(\pi_T)}(P||Q^*) - D_{\sigma(\pi_T)}(\hat{P}_n||\hat{Q}_n^*) \right| \leq r_{b_n, \delta_n}(k) \right\}, \quad (39)$$

for all  $k \in \{1, \dots, |T_{b_n}^{full}|\}$ , we have that the conditions stated on  $(b_n)$  and  $(\delta_n)$  are the weakest to obtain from (15) that  $\lim_{n \rightarrow \infty} \sup_{k \in \{1, \dots, |T_{b_n}^{full}|\}} r_{b_n, \delta_n}(k) = 0$ . This last condition is crucial to being able to apply Lemma 1 in  $\{\mathcal{E}_{\delta_n, b_n}^{n, k}, k\}$ . In fact,  $\mathbb{P}\left(\mathcal{E}_{\delta_n, b_n}^n \equiv \bigcap_{k=1}^{|T_{b_n}^{full}|} \mathcal{E}_{\delta_n, b_n}^{n, k}\right) \geq 1 - \delta_n$  by definition of  $r_{b_n, \delta_n}(k)$  in (16) and the condition expressed in (15).

Importantly under  $\mathcal{H}_0$  (i.e.,  $I(X; Y) = 0$ ), for  $k = 1$ , we can consider that  $r_{b_n, \delta_n}(k) = 0$ , because  $D_{\sigma(\{\mathbb{R}^d\})}(P||Q^*) = D_{\sigma(\{\mathbb{R}^d\})}(\hat{P}_n||\hat{Q}_n^*) = 0$  for any  $z_1^n \in \mathbb{R}^{dn}$ . Consequently, for any  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^n = \bigcap_{k=1}^{|T_{b_n}^{full}|} \mathcal{E}_{\delta_n, b_n}^{n, k}$ , it follows that

$$D_{\sigma(\pi_{\hat{T}_{b_n, \delta_n}})}(\hat{P}_n||\hat{Q}_n^*) - r_{b_n, \delta_n}\left(\left|\hat{T}_{b_n, \delta_n}\right|\right) \geq \sup_{T \in \mathcal{G}_{b_n}^k} D_{\sigma(\pi_T)}(\hat{P}_n||\hat{Q}_n^*) - r_{b_n, \delta_n}(k) \quad (40)$$

$$\geq \sup_{T \in \mathcal{G}_{b_n}^k} D_{\sigma(\pi_T)}(P||Q^*) - 2r_{b_n, \delta_n}(k), \quad (41)$$

for any  $k \in \{1, \dots, |T_{b_n}^{full}|\}$ . The first inequality in (40) is from the definition of  $\hat{T}_{b_n, \delta_n}$  in (17) and the second inequality in (41) is from the fact that if  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^{n, k}$  (see (39)), then

$$\left| \sup_{T \in \mathcal{G}_{b_n}^k} D_{\sigma(\pi_T)}(P||Q^*) - \sup_{T \in \mathcal{G}_{b_n}^k} D_{\sigma(\pi_T)}(\hat{P}_n||\hat{Q}_n^*) \right| \leq r_{b_n, \delta_n}(k). \quad (42)$$

On the other hand, if  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^n$ , then

$$D_{\sigma(\pi_{\hat{T}_{b_n, \delta_n}})}(\hat{P}_n||\hat{Q}_n^*) - r_{b_n, \delta_n}\left(\left|\hat{T}_{b_n, \delta_n}\right|\right) \leq D_{\sigma(\pi_{\hat{T}_{b_n, \delta_n}})}(P||Q^*) \leq D(P||Q^*). \quad (43)$$

<sup>16</sup>This means formally that  $\lim_{n \rightarrow \infty} \frac{\hat{i}_{\pi_{b_n, \delta_n}}(O_1, \dots, O_n)}{n^{-p}} = 0$  with probability one.

At this point, we use the independence assumption<sup>17</sup> and the two inequalities (41) and (43) to obtain that for any  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^n$ :

$$D_{\sigma(\pi_{\hat{T}_{b_n, \delta_n}})}(\hat{P}_n||\hat{Q}_n^*) - r_{b_n, \delta_n}\left(\left|\hat{T}_{b_n, \delta_n}\right|\right) = 0. \quad (44)$$

Finally using that  $r_{b_n, \delta_n}(1) = 0$ , it is simple to verify that for any  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^n$  the trivial tree (with one cell) is a solution for (17), i.e.,  $|\hat{T}_{b_n, \delta_n}| = 1$ . Then from (44),  $D_{\sigma(\pi_{\hat{T}_{b_n, \delta_n}})}(\hat{P}_n||\hat{Q}_n^*) = \hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n)) = 0$ , which concludes the argument using that  $\mathbb{P}(\mathcal{E}_{\delta_n, b_n}^n) \geq 1 - \delta_n$  and the assumption that  $(\delta_n)$  is  $o(1)$ .

Under  $\mathcal{H}_1$  (i.e.,  $I(X; Y) > 0$ ), Silva *et al.* [39, Th. 2, Eq.(33)] showed that for any  $z_1^n \in \mathcal{E}_{\delta_n, b_n}^n$

$$I(X, Y) - \hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n)) \leq \min_{T \ll T_{b_n}^{full}} [I(X; Y) - D_{\sigma(\pi_T)}(P||Q^*(P))] + 2r_{b_n, \delta_n}(T).$$

Again using that  $\mathbb{P}(\mathcal{E}_{\delta_n, b_n}^n) \geq 1 - \delta_n$ , we obtain the bound in (21).  $\square$

#### APPENDIX D PROOF OF LEMMA 3

*Proof:* Let us use the definition of the typical set introduced in (39) and  $\mathcal{E}_{\delta_n, b_n}^n \equiv \bigcap_{k=1}^{|T_{b_n}^{full}|} \mathcal{E}_{\delta_n, b_n}^{n, k}$ . Under the assumption of this result, we know that  $\mathbb{P}(\mathcal{E}_{\delta_n, b_n}^n) \geq 1 - \delta_n$ . In addition, in the proof of Theorem 2, it is shown that if  $z^n \in \mathcal{E}_{\delta_n, b_n}^n$ , then  $|\hat{T}_{b_n, \delta_n}| = 1$ , which means that  $\pi_{b_n, \delta_n} = \{\mathbb{R}^d\}$ . Therefore, we have that<sup>18</sup>

$$\mathcal{E}_{\delta_n, b_n}^n \subset \{z^n \in \mathbb{R}^{dn} : \pi_{b_n, \delta_n}(z^n) = \{\mathbb{R}^d\}\}. \quad (45)$$

Using the definition of  $\mathcal{T}_0((z_n)_{n \geq 1})$ , we have that

$$\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \leq M) = \mathbb{P}_{\mathbf{Z}}\left(\bigcap_{n > M} \{z^n : \pi_{b_n, \delta_n}(z^n) = \{\mathbb{R}^d\}\}\right) \quad (46)$$

$$= 1 - \mathbb{P}_{\mathbf{Z}}\left(\bigcup_{n > M} \{z^n : \pi_{b_n, \delta_n}(z^n) \neq \{\mathbb{R}^d\}\}\right) \quad (47)$$

$$\geq 1 - \sum_{n > M} \mathbb{P}_{\mathbf{Z}}(\{z^n : \pi_{b_n, \delta_n}(z^n) \neq \{\mathbb{R}^d\}\}) \quad (48)$$

$$\geq 1 - \sum_{n > M} \mathbb{P}_{\mathbf{Z}}((\mathcal{E}_{\delta_n, b_n}^n)^c) \geq 1 - \sum_{n > M} \delta_n. \quad (49)$$

The equality in (46) comes from definition in (22). The inequality in (48) comes from the sub-additivity of  $\mathbb{P}_{\mathbf{Z}}$  and the bounds in (49) follow from (45) and the construction of  $\mathcal{E}_{\delta_n, b_n}^n$ . Finally, using the assumption that  $(\delta_n) \in \ell_1(\mathbb{N})$  in (49), we have that

$$\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) < \infty) = \lim_{M \rightarrow \infty} \mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \leq M) = 1. \quad (50)$$

$\square$

<sup>17</sup>Under  $\mathcal{H}_0$ , we have that  $D_{\sigma(\pi_T)}(P||Q^*) = 0$  for any  $T \ll T_{b_n}^{full}$ .

<sup>18</sup>Here we use the notation  $\pi_{b_n, \delta_n}(z^n)$  to make explicit that the partition is data-driven.

APPENDIX E  
PROOF OF THEOREM 3

*Proof:* Under  $\mathcal{H}_1$ , the assumptions on  $(b_n)$  and  $(\delta_n)$  are within the admissible range stated in Theorem 2 part i). Consequently, we have consistency on the regret, i.e.,  $\lim_{n \rightarrow \infty} |i_n(Z_1, \dots, Z_n) - \hat{i}_{\pi_{b_n, \delta_n}}(O_{\pi_{b_n, \delta_n}}(Z_1), \dots, O_{\pi_{b_n, \delta_n}}(Z_n))| = 0$ ,  $\mathbb{P}$ -almost surely. This result is sufficient to obtain that  $(\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n))_{n \geq 1}$  reaches 1 eventually with probability one, from the arguments presented in the proof of Theorem 1.

Under  $\mathcal{H}_0$ , it is useful to define the following last exit time associated with the detection of independence:

$$\begin{aligned} \mathcal{T}((z_n)_{n \geq 1}, (a_n)_{n \geq 1}) &\equiv \\ \sup \{m \geq 1 : \phi_{b_m, \delta_m, a_m}(z_1, \dots, z_m) = 1\} &\in \mathbb{N} \cup \{\infty\}. \end{aligned} \quad (51)$$

Saying that  $(\phi_{b_n, \delta_n, a_n}(Z_1, \dots, Z_n))_{n \geq 1}$  reaches 0 eventually with probability one is equivalent to the condition that

$$\mathbb{P}(\mathcal{T}((Z_n)_{n \geq 1}, (a_n)_{n \geq 1}) < \infty) = 1, \quad (52)$$

from the arguments presented in the proof of Theorem 1. At this point, it is important to revisit the definition of  $\mathcal{T}_0((z_n)_{n \geq 1})$  in (22), where if for some  $z^m \in \mathbb{R}^{dm}$  we have that  $|\hat{T}_{b_m, \delta_m}| = 1$ ; this implies that  $\hat{i}_{\pi_{b_m, \delta_m}}(O_{\pi_{b_m, \delta_m}}(z_1), \dots, O_{\pi_{b_m, \delta_m}}(z_m)) = i_m(z_1, \dots, z_m) = 0$  (details on Appendix C), and therefore,  $\phi_{b_m, \delta_m, a_m}(z_1, \dots, z_m) = 0$  independent of  $a_m$  (see Eq.(7)). Consequently, we have that

$$\begin{aligned} \{z^m \in \mathbb{R}^{dm} : \pi_{b_m, \delta_m}(z^m) = \{\mathbb{R}^d\}\} &= \{z^m : |\hat{T}_{b_m, \delta_m}| = 1\} \\ \subset \{z^m \in \mathbb{R}^{dm} : \phi_{b_m, \delta_m, a_m}(z_1, \dots, z_m) = 0\} \end{aligned} \quad (53)$$

for any  $a_n$ , and it follows that

$$\begin{aligned} \mathbb{P}(\mathcal{T}((Z_n)_{n \geq 1}, (a_n)_{n \geq 1}) \leq M) &= \\ \mathbb{P}_{\mathbf{Z}} \left( \bigcap_{n > M} \{z^n : \phi_{b_n, \delta_n, a_n}(z_1, \dots, z_n) = 0\} \right) \end{aligned} \quad (54)$$

$$\geq \mathbb{P}_{\mathbf{Z}} \left( \bigcap_{n > M} \{z^n : \pi_{b_n, \delta_n}(z^n) = \{\mathbb{R}^d\}\} \right) \quad (55)$$

$$= \mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \leq M). \quad (56)$$

The identity in (54) follows from the definition in (51), and equations (55) and (56) follow from (53) and (46), respectively. The argument concludes from Lemma 3, as we know that under the assumptions stated on  $(b_n)$  and  $(\delta_n)$ ,  $\lim_{M \rightarrow \infty} \mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \leq M) = 1$ , which implies the result in (52) from (56).  $\square$

APPENDIX F  
PROOF OF THEOREM 4

*Proof:* Under  $\mathcal{H}_0$ , we can adopt directly the bound presented in (49):

$$\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) < m) \geq 1 - \sum_{n \geq m} \delta_n = 1 - \sum_{n \geq m} C e^{-n^{1/3}}, \quad (57)$$

where we include the assumption that  $(1/\delta_n)_n \approx (e^{n^{1/3}})$  (i.e.,  $C$  is positive constant). For the series in (57) we have

$$\begin{aligned} \sum_{n \geq m} e^{-n^{1/3}} &= e^{-m^{1/3}} \left[ 1 + \frac{e^{-(m+1)^{1/3}}}{e^{-m^{1/3}}} + \dots + \frac{e^{-(m+j)^{1/3}}}{e^{-m^{1/3}}} + \dots \right] \end{aligned} \quad (58)$$

$$= e^{-m^{1/3}} \cdot \underbrace{\sum_{j \geq 0} e^{-j^{1/3}}}_{\equiv \mathcal{I}}, \quad (59)$$

where  $\mathcal{I} < \infty$ . Therefore, we have that  $\mathbb{P}(\mathcal{T}_0((Z_n)_{n \geq 1}) \geq m) \leq C \mathcal{I} e^{-m^{1/3}}$ .  $\square$

ACKNOWLEDGMENT

The authors thank Diane Greenstein for editing and proof-reading all this material.

REFERENCES

- [1] C.-J. Ku and T. L. Fine, "Testing for stochastic independence: Application to blind source separation," *IEEE Trans. Signal Process.*, vol. 53, no. 5, pp. 1815–1826, May 2005.
- [2] P. Comon, "Independent component analysis: A new concept?," *Signal Process.*, vol. 36, pp. 287–314, 1994.
- [3] J. F. Cardoso, "Blind separation of instantaneous mixtures of non-stationary sources," *IEEE Trans. Signal Process.*, vol. 49, no. 9, pp. 1837–1848, Sep. 2001.
- [4] J. C. Lv, K. K. Tan, Z. Yi, and S. Huang, "Convergence analysis of a class of HyvÄrinen-Oja's ICA learning algorithms with constant learning rates," *IEEE Trans. Signal Process.*, vol. 57, no. 5, pp. 1811–1824, May 2009.
- [5] B. Bloem-Reddy and Y. W. Teh, "Probabilistic symmetries and invariant neural networks," *J. Mach. Learn. Res.*, vol. 20, pp. 1–61, 2020. [Online]. Available: <https://arxiv.org/abs/1901.06082>
- [6] J. Range, "Conditional independence testing based on a nearest-neighbor estimator of conditional mutual information," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2018, pp. 938–947.
- [7] A. Bellot and M. V. D. Schaar, "Conditional independence testing using generative adversarial networks," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 2202–2211.
- [8] R. Sen, A. T. Suresh, K. Shanmugam, A. G. Dimakis, and S. Shakkottai, "Model-powered conditional independence test," in *Proc. Conf. Neural Inf. Process. Syst.*, 2017, pp. 2955–2965.
- [9] C. Ku and T. Fine, "A Bayesian independence test for small datasets," *IEEE Trans. Signal Process.*, vol. 54, no. 10, pp. 4026–4031, Oct. 2006.
- [10] D. Reshef *et al.*, "Detecting novel associations in large data sets," *Science*, vol. 334, no. 6062, pp. 1518–1524, 2011.
- [11] A. Gretton, O. Busquet, A. Smola, and B. Scholkopf, "Measuring statistical dependence with Hilbert-Schmit norms," in *Proc. 16th Int. Conf. Algorithms Learn. Theory*, 2005, pp. 63–78.
- [12] A. Gretton and L. Györfi, "Nonparametric independence tests: Space partitioning and kernel approaches," in *Proc. 19th Int. Conf. Algorithms Learn. Theory*, Springer, 2008, pp. 183–198.
- [13] A. Gretton, K. Fukumizu, C. H. Teo, L. Song, B. Schölkopf, and A. J. Smola, "A kernel statistical test of independence," in *Proc. Adv. Neural Inf. Process. Syst.*, 2007, pp. 585–592.

- [14] N. Pfister, P. Bühlmann, B. Schölkopf, and J. Peters, “Kernel-based tests for joint independence,” *J. Roy. Statist. Soc.: Ser. B (Statist. Methodology)*, vol. 80, no. 1, pp. 5–31, 2018.
- [15] G. Székely, M. Rizzo, and N. Bakirov, “Measuring and testing dependency by correlation of distances,” *Ann. Statist.*, vol. 35, no. 6, pp. 2769–2794, 2007.
- [16] G. J. Székely and M. L. Rizzo, “Brownian distance covariance,” *Ann. Appl. Statist.*, vol. 3, pp. 1236–1265, 2009.
- [17] G. J. Székely, “The distance correlation t-test of independence in high dimension,” *J. Multivariate Anal.*, vol. 117, pp. 193–213, 2013.
- [18] S. Zheng, N. Shi, and Z. Zhang, “Generalized measures of correlation for asymmetry, nonlinearity, and beyond,” *J. Amer. Stat. Assoc.*, vol. 107, pp. 1239–1252, 2012.
- [19] X. Wang, B. Jiang, and J. Liu, “Generalized R-squared for detecting non-independence,” *Biometrika*, vol. 104, pp. 129–139, 2017.
- [20] A. Gretton and L. Györfi, “Consistent nonparametric tests of independence,” *J. Mach. Learning Res.*, vol. 11, pp. 1391–1423, 2010.
- [21] L. Ma and J. Mao, “Fisher exact scanning for dependency,” *J. Amer. Stat. Assoc.*, vol. 114, no. 525, pp. 245–258, 2019.
- [22] K. Zhang, “Bet on independence,” *J. Amer. Stat. Assoc.*, vol. 114, no. 528, pp. 1620–1637, 2019.
- [23] R. Heller, Y. Heller, S. Kaufman, B. Brill, and M. Gorfine, “Consistent distribution-free K-sample and independence tests for univariate random variables,” *J. Mach. Learn. Res.*, vol. 17, pp. 1–54, 2016.
- [24] R. Heller and Y. Heller, “Multivariate tests of association based on univariate tests,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, vol. 29, pp. 208–216.
- [25] M. Menedez, J. Pardo, L. Pardo, and K. Zografos, “On tests of independence based on minimum  $\phi$ -divergence estimator with constraint: An application to modeling DNA,” *Comput. Statist. Data Anal.*, vol. 51, pp. 1100–1118, 2006.
- [26] F. Han, S. Chen, and H. Liu, “Distribution-free tests of independence in high dimensions,” *Biometrika*, vol. 104, pp. 813–828, 2017.
- [27] J. Dauxois and G. M. Nkiet, “Nonlinear canonical analysis and independence tests,” *Ann. Statist.*, vol. 26, no. 4, pp. 1254–1278, 1998.
- [28] A. Dionisio, R. Menezes, and D. A. Mendes, “Entropy-based independence test,” *Nonlinear Dyn.*, vol. 44, pp. 351–357, 2006.
- [29] J. Suzuki, “An estimator of mutual information and its application to independence testing,” *Entropy*, vol. 18, no. 109, pp. 1–18, Mar. 2016.
- [30] I. Kontoyiannis and M. Skoularidou, “Estimating the directed information and testing for causality,” *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6053–6067, Nov. 2016.
- [31] L. Devroye and G. Lugosi, *Combinatorial Methods in Density Estimation*. New York, NY, USA: Springer, 2001.
- [32] S. Kullback, *Information Theory and Statistics*. New York, NY, USA: Wiley, 1958.
- [33] H. Shi, M. Drton, and F. Han, “Distribution-free consistent independence tests via center-outward ranks and signs,” *J. Amer. Stat. Assoc.*, pp. 1–16, 2020.
- [34] K. Zhang, Z. Zhao, and W. Zhou, “Beauty powered beast,” 2021, [arXiv:2103.00674](https://arxiv.org/abs/2103.00674).
- [35] I. Csizsar and P. Shields, *Information Theory and Statistics: A Tutorial*. Boston, MA, USA: Now, 2004.
- [36] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., New York, NY, USA: Wiley, 2006.
- [37] G. Lugosi and A. B. Nobel, “Consistency of data-driven histogram methods for density estimation and classification,” *Ann. Statist.*, vol. 24, no. 2, pp. 687–706, 1996.
- [38] G. A. Darbellay and I. Vajda, “Estimation of the information by an adaptive partition of the observation space,” *IEEE Trans. Inf. Theory*, vol. 45, no. 4, pp. 1315–1321, May 1999.
- [39] J. F. Silva and S. Narayanan, “Complexity-regularized tree-structured partition for mutual information estimation,” *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1940–1952, Mar. 2012.
- [40] J. F. Silva and S. Narayanan, “Non-product data-dependent partitions for mutual information estimation: Strong consistency and applications,” *IEEE Trans. Signal Process.*, vol. 58, no. 7, pp. 3497–3511, Jul. 2010.
- [41] J. F. Silva and S. Narayanan, “Information divergence estimation based on data-dependent partitions,” *J. Stat. Plan. Inference*, vol. 140, no. 11, pp. 3180–3198, Nov. 2010.
- [42] R. E. Blahut, “Hypothesis testing and information theory,” *IEEE Trans. Inf. Theory*, vol. 20, no. 4, pp. 405–417, Jul. 1974.
- [43] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*. Belmont, CA, USA: Wadsworth, 1984.
- [44] C. Scott and R. D. Nowak, “Minimax-optimal classification with dyadic decision trees,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1335–1353, Apr. 2006.
- [45] C. Scott, “Tree pruning with subadditive penalties,” *IEEE Trans. Signal Process.*, vol. 53, no. 12, pp. 4518–4525, Dec. 2005.
- [46] L. Devroye, L. Györfi, and G. Lugosi, *A Probabilistic Theory of Pattern Recognition*. New York, NY, USA: Springer, 1996.
- [47] M. Gonzalez and J. F. Silva, “Data-driven representations for testing independence: A connection with mutual information estimation,” in *Proc. IEEE Int. Symp. Inf. Theory*, 2020, pp. 1–5.
- [48] J. F. Silva and S. Narayanan, “Universal consistency of data-driven partitions for divergence estimation,” in *Proc. IEEE Int. Symp. Inf. Theory*, 2007, pp. 2021–2025.
- [49] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. Cambridge, MA, USA: MIT Press, 2009.
- [50] P. J. Bickel, C. A. J. Klaassen, Y. Ritov, and J. A. Wellner, *Efficient and Adaptive Estimation for Semiparametric Models*. New York, NY, USA: Springer, 1998.
- [51] N. Merhav and M. Feder, “Universal prediction,” *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2124–2147, Oct. 1998.
- [52] D. Bontemps, S. Boucheron, and E. Gassiat, “About adaptive coding on countable alphabets,” *IEEE Trans. Inf. Theory*, vol. 60, no. 2, pp. 808–821, Feb. 2014.

**Mauricio E. Gonzalez** received the bachelor’s and M.Sc. degrees in electrical engineering from the University of Chile, Santiago, Chile, in 2019 and 2021, respectively. He was a Research Assistant with the Information and Decision System Group during 2019–2021 and was a Teaching Assistant for many undergraduate and graduate courses with the University of Chile. His research interests include probabilities, statistics, control systems, signal processing, information theory, machine learning and their application to the study of failure prognostics and tests of independence.

**Jorge F. Silva** (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in electrical engineering from the University of Southern California (USC), Los Angeles, CA, USA, 2005 and 2008, respectively. He is currently an Associate Professor with the Department of Electrical Engineering, University of Chile, Santiago, Chile, and the Principal Investigator with the Advanced Center for Electrical and Electronic Engineering (UTFSM). Jorge F. Silva was a Research Assistant with the Signal Analysis and Interpretation Laboratory (SAIL), USC, during 2003–2008 and was also a Research Intern with the Speech Research Group, Microsoft Corporation, Redmond, in summer 2005. His research interests include information measure estimation, machine learning, representation learning, statistical learning, universal source coding, compressed sensing, sparse processes, wavelet, and multiresolution analysis. He was the recipient of the Outstanding Thesis Award 2009 for Theoretical Research of the Viterbi School of Engineering, the Viterbi Doctoral Fellowship 2007–2008 and Simon Ramo Scholarship 2007–2008 USC. He was an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING February 2006–February 2008.

**Miguel Videla** received the bachelor’s and M.Sc. degrees in electrical engineering from the University of Chile, Santiago, Chile, in 2019 and 2021, respectively. He was a Research Assistant with the Information and Decision System Group during 2019–2021 and with the French-Chilean Laboratory for Astronomy in 2021. His research interests include probabilities, statistics, machine learning, statistical learning, information theory, and astrostatistics.

**Marcos E. Orchard** (Member, IEEE) received the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2005 and 2007, respectively. He is currently a Professor with the Department of Electrical Engineering, Universidad de Chile, Santiago, Chile, and an Associate Researcher with the Advanced Center for Electrical and Electronic Engineering (UTFSM). He has authored or coauthored more than 100 articles. His current research interests include failure prognosis, prognostics decision-making, electromobility, mining industry, and finance. He is a Fellow of the Prognostic and Health Management Society (PHMS) and the Editor-in-Chief of the *International Journal of Prognostics and Health Management*.