An Inferential Measure of Dependence Between Two Systems Using Bayesian Model Comparison

Guillaume Marrelec^D and Alain Giron^D

Abstract-We propose to quantify dependence between two systems \mathcal{X} and \mathcal{Y} in a dataset D based on the Bayesian comparison of two models: one, H_0 , of statistical independence and another one, H_1 , of dependence. In this framework, dependence between \mathcal{X} and \mathcal{Y} in D, denoted $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$, is quantified as $P(H_1|D)$, the posterior probability for the model of dependence given D, or any strictly increasing function thereof. It is therefore a measure of the evidence for dependence between ${\cal X}$ and \mathcal{Y} as modeled by H_1 and observed in D. We review several statistical models and reconsider standard results in the light of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ as a measure of dependence. Using simulations, we focus on two specific issues: 1) the effect of noise and 2) the behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ when H_1 has a parameter coding for the intensity of dependence. We then derive some general properties of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$, showing that it quantifies the information contained in D in favor of H_1 versus H_0 . While some of these properties are typical of what is expected from a valid measure of dependence, others are novel and naturally appear as desired features for specific measures of dependence, which we call *inferential*. We finally put these results in perspective; in particular, we discuss the consequences of using the Bayesian framework as well as the similarities and differences between $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ and mutual information.

Index Terms-Bayesian model comparison, dependence, independence, measure of dependence.

I. INTRODUCTION

NDEPENDENCE and dependence are key concepts in science whose goal is to characterize the structural relationships between systems. Consider two systems $\mathcal X$ and \mathcal{Y} characterized by a probabilistic description in terms of (possibly multivariate) random variables X and Y, respectively, with known joint probability distribution $f_{XY}(\mathbf{x}, \mathbf{y})$ and marginals $f_X(\mathbf{x})$ and $f_Y(\mathbf{y})$. In this context, \mathcal{X} and \mathcal{Y} are said to be independent if they underlying random variables are independent, i.e., [1, Sec. 2.2]

$$f_{XY}(\boldsymbol{x}, \boldsymbol{y}) = f_X(\boldsymbol{x}) f_Y(\boldsymbol{y}). \tag{1}$$

When X and Y are not independent, $f_{XY}(\mathbf{x}, \mathbf{y})$ differ from $f_X(\mathbf{x}) f_Y(\mathbf{y})$, and both systems are said to be dependent. In this

Received 3 January 2024; revised 12 April 2024 and 8 July 2024; accepted 20 November 2024. Date of publication 17 December 2024; date of current version 19 February 2025. This article was recommended by Associate Editor Z. Wei. (Corresponding author: Guillaume Marrelec.)

The authors are with the Laboratoire d'Imagerie Biomédicale, Sorbonne Universite, CNRS, INSERM, 75006 Paris, France (e-mail: guillaume.marrelec@inserm.fr).

This article has supplementary material provided by the authors and color versions of one or more figures available at https://doi.org/10.1109/TSMC.2024.3509025.

Digital Object Identifier 10.1109/TSMC.2024.3509025

case, an important issue is the quantification of dependence between \mathcal{X} and \mathcal{Y} , where one tries to measure to what degree $f_{XY}(\mathbf{x}, \mathbf{y})$ differs from $f_X(\mathbf{x}) f_Y(\mathbf{y})$. There are many ways to depart from independence, and this question remains open in the general case. Many measures of dependence have been proposed, including but not limited to mutual information [2, Secs. 2 and 8], maximal correlation coefficient [3], [4], the mixed derivative measure of marginal interaction [1, Sec. 2.3], Hoeffding's procedure [5], distance correlation [6], [7], circular correlation [8], and Hilbert-Schmidt information criterion [9].

Once a theoretical measure has been proposed, another layer of complexity is often added by the fact that we do not know $f_{XY}(x, y)$. A common situation is when one knows (or assumes) that it belongs to a family $f_{XY}(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta})$ parameterized by an (unknown) parameter θ which has to be estimated using a dataset of N independent and identically distributed (i.i.d.) samples (x_n, y_n) . Various estimation strategies have been proposed, from plug-in estimators (where one computes the measure by replacing the true value of the parameter by its estimator) [10, Ch. 9], [11, Ch. 12, Sec. 3.6] to more refined techniques proposing estimators for the measure itself [12], [13], [14]. The interest of this kind of approaches relies on the asymptotic convergence of the estimators toward the true value of the measure. However, the values taken by these estimators for finite data size do not have simple interpretations.

In this article, we introduce another general measure of dependence $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ between systems \mathcal{X} and \mathcal{Y} given dataset D. Our starting point is the real-life situation where the joint behavior of $(\mathcal{X}, \mathcal{Y})$ is characterized by a dataset of N i.i.d. samples (x_n, y_n) . We propose to characterize dependence between \mathcal{X} and \mathcal{Y} in D by using a Bayesian inference scheme that compares the credibility of two competing models: 1) H_0 , which describes \mathcal{X} and \mathcal{Y} as independent and 2) H_1 , which describes them as dependent. This setting was already used to infer the structure of independence within a multivariate distribution [15], [16]. In the case of bivariate discrete distributions and multivariate normal distributions, a connection was also observed between the log posterior odd ratio and mutual information [15], [16], [17]. We here propose to go one step further and advocate that the Bayesian comparison of models H_1 and H_0 mentioned above provides a family of measures that can be used to quantify the level of dependence between \mathcal{X} and \mathcal{Y} . As will be further developed below, any measure $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ in this family is an increasing function of correlation, mutual information, the minimum discrimination

2168-2216 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

information, and the log-likelihood ratio criterion for testing independence, in cases where it makes sense to define such measures. As a consequence, $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ shares many features with these classical measures of dependence. Unlike these same measures, though, the value of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ for a given D and finite N has a direct, exact interpretation as (a strictly increasing function of) the probability that X and Ybe dependent (as described by H_1) for D. In other words, $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ quantifies the evidence for—or credibility of dependence between \mathcal{X} and \mathcal{Y} as modeled by H_1 and observed in D, i.e., the amount of information contained in D in favor of H_1 versus H_0 . It has several interesting features regarding the effect of data size and noise, which are detailed throughout the manuscript and characterize what we would call an *inferential* measure of dependence.

The outline of our manuscript is the following. In Section II, we present the general theoretical framework. In Section III, we use simulation studies to investigate two specific issues: 1) the effect of noise and 2) the behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ when H_1 has a parameter coding for the intensity of dependence. Section IV provides a real-life application in neuroscience and neuroimaging. We then state some key properties of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ in Section V. Further issues are discussed in Section VI.

II. BAYESIAN MEASURES OF DEPENDENCE

In this section, we present the core of our method. After a quick review of Bayesian model comparison for the investigation of statistical independence and dependence (Section II-A), we introduce a general measure of dependence $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ which can either take the form of a posterior probability or any strictly increasing function thereof (Section II-B). We then investigate the theoretical properties of $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$, a particular instance of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$, in the case of known distributions (Section II-C) and known likelihood functions with unknown parameters (Section II-D). In Section II-E, we apply our framework to the important case where H_0 is nested in H_1 , a framework that can be applied to several common models, such as maximum-entropy distributions, multivariate normal distributions, and bivariate discrete distributions. In Section II-F, we consider the other particular case where dependence is modeled through a copula. Section II-G investigates the consequences of model misspecification. Finally, in Section II-H, we summarize the main results obtained thus far.

A. Model Comparison

We here provide a quick description of the general framework of Bayesian model comparison for the investigation of statistical dependence. For more details, the reader can refer to [15] or [16]. Consider the following two competing models.

1) A model H_0 (with parameter $\boldsymbol{\theta}^{(0)} \in \Theta^{(0)}$) in which \mathcal{X} and \mathcal{Y} are independent and where the likelihood is given by

$$l_0\left(\boldsymbol{x}, \boldsymbol{y}|\boldsymbol{\theta}^{(0)}\right) = f_X^{(0)}\left(\boldsymbol{x}|\boldsymbol{\theta}^{(0)}\right) f_Y^{(0)}\left(\boldsymbol{y}|\boldsymbol{\theta}^{(0)}\right).$$
(2)

2) Another model H_1 (with model parameter $\theta^{(1)} \in \Theta^{(1)}$) in which \mathcal{X} and \mathcal{Y} are dependent and the likelihood is

$$l_1(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}^{(1)}) = f_{XY}^{(1)}(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}^{(1)}).$$
(3)

In a Bayesian framework, all the information from the data that is relevant for the problem at hand is summarized by the posterior probabilities of H_0 and H_1 given the dataset D, $p(H_0|D)$ and $p(H_1|D)$, respectively. Both quantities can be calculated using Bayes updating rule

$$p(H_i|D) = \frac{p(H_i) p(D|H_i)}{p(D)}, \qquad i = 0, 1.$$
 (4)

 $p(H_i)$ is the model prior probability and can be set depending on the prior belief that we have regarding the relative plausibility of both competing hypotheses. As to the marginal model likelihood $p(D|H_i)$, it can be expressed using the marginalization formula

$$p(D|H_i) = \int_{\boldsymbol{\theta}^{(i)} \in \Theta^{(i)}} p\left(\boldsymbol{\theta}^{(i)}|H_i\right) p\left(D|H_i, \boldsymbol{\theta}^{(i)}\right) d\boldsymbol{\theta}^{(i)}.$$
 (5)

In this expression, $p(\theta^{(i)}|H_i)$ of is the parameter prior. We set it to $h_i(\theta^{(i)})$ for $\theta^{(i)} \in \Theta^{(i)}$. $p(D|H_i, \theta^{(i)})$ is the likelihood function. For *N* i.i.d. samples

$$D = \left\{ \left(\boldsymbol{x}_1, \boldsymbol{y}_1 \right), \dots, \left(\boldsymbol{x}_N, \boldsymbol{y}_N \right) \right\}$$
(6)

it can be decomposed as

$$p(D|H_i, \boldsymbol{\theta}^{(i)}) = \prod_{n=1}^N l_i(\boldsymbol{x}_n, \boldsymbol{y}_n | \boldsymbol{\theta}^{(i)}).$$
(7)

B. A Family of Measures

By definition, $p(H_1|D)$ quantifies the posterior probability that \mathcal{X} and \mathcal{Y} be dependent. As a consequence, it can be considered as a measure of dependence between \mathcal{X} and \mathcal{Y}

$$\mathfrak{B}_{\rm pr}(\mathcal{X}, \mathcal{Y}|D) \equiv p(H_1|D). \tag{8}$$

Indeed, it is equal to 0 when it is known that \mathcal{X} and \mathcal{Y} are independent as in H_0 , equal to 1 when they are known to be dependent in agreement with H_1 , and in between when we are sure of neither; the larger it is, the more probable the dependence of \mathcal{X} and \mathcal{Y} is.

If one agrees to treat $p(H_1|D)$ as a measure of dependence, then various strictly increasing mappings of it could also be considered, with different ranges. For instance, one could use the posterior odd ratio of H_1 versus H_0

$$\mathfrak{B}_{\rm r}(\mathcal{X},\mathcal{Y}|D) \equiv \frac{p(H_1|D)}{p(H_0|D)} = \frac{p(H_1)}{p(H_0)} \frac{p(D|H_1)}{p(D|H_0)}.$$
 (9)

The first fraction of the right-hand side equation is the socalled prior odd ratio, while the second fraction is the Bayes factor. The Bayes factor itself could be used as a measure of dependence. We can also use the log scale to obtain a better representation of the measure of dependence, either in its usual form

$$\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D) \equiv \ln \frac{p(H_1)}{p(H_0)} + \ln \frac{p(D|H_1)}{p(D|H_0)}$$
(10)

or in log₁₀, so that a value of *b* means that H_1 is 10^{*b*} times more probable than H_0 ; or expressed, such as to yield values on a scale similar to decibels [18]. Measures similar to $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ and $\mathfrak{B}_{dB}(\mathcal{X}, \mathcal{Y}|D)$ but relying on $\mathfrak{B}_{BF}(\mathcal{X}, \mathcal{Y}|D)$ instead of $\mathfrak{B}_r(\mathcal{X}, \mathcal{Y}|D)$ could be proposed (i.e., not taking the relative priors of H_0 and H_1 into account).

Importantly, all these measures are strictly increasing function of dependence, i.e., they increase as the probability for H_1 (dependence) increases.

In the following, we will mostly focus on $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$ and $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$, as it is these measures whose properties are easiest to investigate. Furthermore, they show the closest connections to existing frameworks for quantifying dependence, such as the log-likelihood criterion for independence and mutual information.

C. Known Distributions

If we assume that the distribution of (X, Y) is known exactly in both H_0 and H_1 , then it is straightforward to show that (see Section 1 of Supplementary Material)

- If H₀ is true (independence), 𝔅_{lnr}(𝔅, 𝔅|D) tends to −∞ as N → ∞, and its sampling expectation is a decreasing function of N.
- 2) If H_1 is true (dependence), $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$ tends to $+\infty$ as $N \to \infty$, and its expectation is an increasing function of *N* and *I*(*X*, *Y*), the mutual information between *X* and *Y*.

The rest of the section is devoted to show that these results also hold for more general models, albeit in a weaker, asymptotic form.

D. Known Likelihood Functions With Unknown Parameters

We now consider the more general case where each likelihood function is not exactly known but belongs to a known family with unknown parameter. While an exact expression for $\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)$ cannot be obtained in this case, we can still derive an asymptotic approximation in a fashion similar to [19, Secs. 7.22–7.27], which itself relies on the consistency of the maximum-likelihood estimate.

For $i \in \{0, 1\}$, let $L_i(\theta^{(i)}) = p(D|H_i, \theta^{(i)})$ be the likelihood function associated with model H_i , as defined in (7). The marginal model likelihoods of (5) can formally be expressed as

$$P(D|H_i) = \int_{\boldsymbol{\theta}^{(i)} \in \Theta^{(i)}} h_i(\boldsymbol{\theta}^{(i)}) L_i(\boldsymbol{\theta}^{(i)}) d\boldsymbol{\theta}^{(i)}.$$
(11)

We assume that $L_i(\theta^{(i)})$ is unimodal and set $\hat{\theta}_N^{(i)}$ its maximumlikelihood estimate

$$\hat{\boldsymbol{\theta}}_{N}^{(i)} = \operatorname{argmax}_{\boldsymbol{\theta}^{(i)} \in \Theta^{(i)}} L_{i} \left(\boldsymbol{\theta}^{(i)} \right), \qquad i = 0, 1.$$
(12)

Assuming that the prior $h_i(\theta^{(i)})$ is strictly positive and of slow variation around $\hat{\theta}_N^{(i)}$, the integral of (11) can be approximated using Laplace method [20], [21] (see Section 2 of

Supplementary Material), yielding

$$\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D) = \ln \frac{L_1(\hat{\boldsymbol{\theta}}_N^{(1)})}{L_0(\hat{\boldsymbol{\theta}}_N^{(0)})} - \frac{D_2 - D_1}{2} \ln N + O(1)$$
(13)

and

$$\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D) = N \Biggl[\hat{I}(X, Y) + \frac{1}{N} \sum_{n=1}^{N} \ln \frac{f_X^{(1)}(\mathbf{x}_n | \hat{\boldsymbol{\theta}}_N^{(1)})}{f_X^{(0)}(\mathbf{x}_n | \hat{\boldsymbol{\theta}}_N^{(0)})} + \frac{1}{N} \sum_{n=1}^{N} \ln \frac{f_Y^{(1)}(\mathbf{y}_n | \hat{\boldsymbol{\theta}}_N^{(1)})}{f_Y^{(0)}(\mathbf{y}_n | \hat{\boldsymbol{\theta}}_N^{(0)})} \Biggr] - \frac{D_2 - D_1}{2} \ln N + O(1).$$
(14)

In (13), the first term of the right-hand side is the classical likelihood ratio test statistic, showing the connection between our approach and the log-likelihood ratio criterion for testing independence [10, Sec. 9.2]. The second term of the right-hand side is the BIC correction [22]. In (14), $\hat{I}(X, Y)$ is the sampling mutual information under H_1 .

In the absence of further assumption, nothing can be said about the asymptotic behavior of this quantity, which depends on the true likelihood function as well as the limits of the maximum-likelihood estimates $\hat{\theta}_N^{(0)}$ and $\hat{\theta}_N^{(1)}$. These limits, in turn, are not necessary well defined in the general case.

E. Nested Models

We here consider the particular case where H_0 is nested in H_1 . In this case, the models allowed by H_0 are included in those allowed by H_1 , i.e., models of H_0 are particular cases of models in H_1 . In other words, there exists a function π such that, for any $\boldsymbol{\theta}^{(0)} \in \Theta^{(0)}$, we have

$$l_0(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta}^{(0)}) = l_1[\mathbf{x}, \mathbf{y}|\pi\left(\boldsymbol{\theta}^{(0)}\right)].$$
(15)

For the sake of simplicity, it is often assumed that $\pi(\boldsymbol{\theta}^{(0)})$ is a projection, i.e., $\Theta^{(1)}$ can be parameterized by $\boldsymbol{\theta}^{(1)} = (\boldsymbol{\theta}^{(0)}, \boldsymbol{\phi})$ with $\boldsymbol{\theta}^{(0)} \in \Theta^{(0)}$ such that $\pi(\boldsymbol{\theta}^{(0)}) = (\boldsymbol{\theta}^{(0)}, \boldsymbol{\phi}_0)$ (i.e., $\boldsymbol{\phi} = \boldsymbol{\phi}_0$). Then it can be shown the following results (see Section 3 of Supplementary Material).

a) Under H_0 : When $N \to \infty$, $\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)$ essentially behaves as $-(1/2)(D_1 - D_0) \ln N$, which is a decreasing function of N that tends to $-\infty$. Also, $\mathbb{E}[\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)|H_0]$ is a decreasing function of N.

b) Under H_1 : When $N \to \infty$, $\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)$ is approximately linearly increasing in N. $\mathbb{E}[\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)|H_1]$ has a first-order approximation that is an increasing function of N, but a second-order approximation that may first decrease before it increases.

Nested models in the particular cases of maximum-entropy distributions, multivariate normal distributions, and bivariate discrete distributions are considered in Sections 4–6 of Supplementary Material, respectively, together with details about the connection between our method and existing methods, such as the log-likelihood ratio criterion, the BIC, mutual information, and the minimum discrimination information statistic.

F. Copula Models

Another particular case of interest is when H_0 and H_1 share the same assumptions regarding marginals for X and Y—modeled by $f_X(\mathbf{x}|\boldsymbol{\phi})$ and $f_Y(\mathbf{y}|\boldsymbol{\phi})$, respectively, and H_1 models dependence through a copula with density $c(\mathbf{u}, \mathbf{v}|\boldsymbol{\psi})$ [23, Ch. 8]. The model parameters are therefore $\boldsymbol{\theta}^{(0)} = \boldsymbol{\phi}^{(0)}$ for H_0 and $\boldsymbol{\theta}^{(1)} = (\boldsymbol{\phi}^{(1)}, \boldsymbol{\psi}^{(1)})$ for H_1 . A common approach to estimate parameters is the method of inference function for margins (IFM) [24, Sec. 10.1], which has a simple interpretation in our framework: first estimate $\boldsymbol{\phi}^{(0)}$ using the maximum-likelihood estimate $\hat{\boldsymbol{\theta}}_N^{(0)} = \hat{\boldsymbol{\phi}}_N^{(0)}$, and then estimate $\boldsymbol{\theta}^{(1)}$ by first setting $\hat{\boldsymbol{\phi}}_N^{(1)} = \hat{\boldsymbol{\phi}}_N^{(0)}$ and then finding the maximumlikelihood estimate $\hat{\boldsymbol{\theta}}_N^{(1)} = (\hat{\boldsymbol{\phi}}_N^{(0)}, \hat{\boldsymbol{\psi}}_N^{(1)})$. In this case, (14) simplifies to

$$\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D) = N\hat{I}(X, Y) - \frac{D_2 - D_1}{2}\ln N + O(1)$$
 (16)

with

$$\hat{I}(X,Y) = \frac{1}{N} \sum_{n=1}^{N} \ln c \Big[F_X \Big(\mathbf{x}_n | \hat{\boldsymbol{\phi}}_N^{(0)} \Big), F_Y \Big(\mathbf{y}_n | \hat{\boldsymbol{\phi}}_N^{(0)} \Big) \Big| \hat{\boldsymbol{\psi}}_N^{(1)} \Big]$$
(17)

where $F_X(\mathbf{x}|\boldsymbol{\phi})$ and $F_Y(\mathbf{y}|\boldsymbol{\phi})$ are the cumulative distribution functions of x and Y, respectively. This quantity is related to minus the entropy of $c(\mathbf{u}, \mathbf{v}|\boldsymbol{\psi})$ [25].

G. Model Misspecification

We now investigate the consequence of considering a true underlying generative model that is neither H_0 nor H_1 . To this aim, we assume that the generative distribution is $f(\mathbf{x}, \mathbf{y}|\boldsymbol{\theta})$, that the estimators of the model parameters under H_0 and H_1 have limits, i.e., $\hat{\boldsymbol{\theta}}_N^{(0)} \xrightarrow{N \to \infty} \boldsymbol{\theta}_\infty^{(0)}$ and $\hat{\boldsymbol{\theta}}_N^{(1)} \xrightarrow{N \to \infty} \boldsymbol{\theta}_\infty^{(1)}$, and that these limits are such that Laplace approximation which can be applied. In (13), the leading term of $\mathcal{X}, \mathcal{Y}|D$ is $\ln [L_1(\hat{\boldsymbol{\theta}}_N^{(1)})/L_0(\hat{\boldsymbol{\theta}}_N^{(0)})]$ can be expressed as (see Section 7 of Supplementary Material)

$$N\left\{D_{\mathrm{KL}}\left[f(\boldsymbol{x},\boldsymbol{y}|\boldsymbol{\theta})\|f_{X}^{(0)}\left(\boldsymbol{x}_{n}|\boldsymbol{\theta}_{\infty}^{(0)}\right)f_{Y}^{(0)}\left(\boldsymbol{y}_{n}|\boldsymbol{\theta}_{\infty}^{(0)}\right)\right] - D_{\mathrm{KL}}\left[f(\boldsymbol{x},\boldsymbol{y}|\boldsymbol{\theta})\|f_{XY}^{(1)}\left(\boldsymbol{x},\boldsymbol{y}|\boldsymbol{\theta}_{\infty}^{(1)}\right)\right]\right\} + o(N) \quad (18)$$

where $o(\cdot)$ is the usual little-o notation. If both Kullback– Leibler divergences in this equation differ, then the likelihood ratio is roughly linear in N, with a proportionality factor whose sign is given by the difference in Kullback–Leibler divergences. As a consequence, for N large enough, if the model corresponding to H_0 is closer to the true generative model (as measured by Kullback–Leibler divergence), $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$ will be decreasing and will tend to $-\infty$; if it is the model corresponding to H_1 that is closer, $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$ will be increasing and will tend to $+\infty$.

H. Summary

In Section II, we introduced a general measure $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ to quantify statistical dependence in data sets. Our framework is based on the Bayesian comparison of a model H_1 taking dependence into account (in a form specified by the model)

and a model H_0 not taking this dependence into account (Section II-A). We then defined $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ as $p(H_1|D)$ or any strictly increasing functions thereof (Section II-B). We investigated the behavior of these measures on i.i.d. data in the case of a known distribution (Section II-C) or a known likelihood family with unknown parameters (Section II-D). We then delved into two particular cases: 1) nested models (Section II-E), which can be applied to maximum-entropy distributions, multivariate normal distributions, and bivariate discrete distributions and 2) copula models of dependence (Section II-F). We finally considered the consequences of model misspecification (Section II-G). In all the cases, we showed that $\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)$ asymptotically behaved as follows.

- 1) Under H_0 , $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D)$ is a decreasing function of N which tends to $-\infty$ as $N \to \infty$;
- Under H₁, 𝔅_{Inr}(𝔅, 𝔅|D) is an increasing function of N (possibly after an initial decrease) which tends to +∞ as N → ∞.
- 3) $\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)$ is an increasing function of $\hat{I}(X, Y)$ which tends to $+\infty$ as $I(X, Y) \to +\infty$.
- 4) In case of model misspecification, $\mathfrak{B}_{\ln r}(\mathcal{X}, \mathcal{Y}|D)$ behaves as if the model closer to the true generative one (in terms of Kullback–Leibler divergence) were the true one.

While these results involve developments that are standard in statistical theory and information theory, they shed some important light on the relevance of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ as a valid measure of dependence. We will come back to this point in Section V.

III. SIMULATION STUDY

In the previous section, we provided general results regarding some common statistical models of dependence where we had direct access to the variables of interest. Here, we use synthetic data to focus on two specific issues: 1) the effect of noise and 2) the behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ when H_1 has a parameter coding for the intensity of dependence.

To investigate the effect of noise, we considered synthetic data originating from three distinct models: two variables following a bivariate normal distribution plus noise (Section III-A), two variables related by a functional relationship plus noise (Section III-B), and two chaotic systems (Section III-D). In all three examples, we varied the size of the dataset N and the variance of the noise σ^2 . We predicted that a good measure of dependence should behave as follows: as the dataset becomes more and more informative (i.e., as N increases and σ^2 decreases), the value of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ should 1) decrease and get increasingly closer to its lower bound if the true underlying model is a model of independence (H_0) and 2) increase and get increasingly closer to its upper bound if the true underlying model is a model of dependence (H_1) .

The behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ when H_1 has a parameter coding for the intensity of dependence was investigated using three models as well: 1) the abovementioned model



Fig. 1. Simulation study: bivariate normal distribution with noise. Boxplots (median and [25%, 75%] percentile) of $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ in various conditions. Top left: Effect of *N* and σ^2 for simulations with $\rho = 0$. Top right: Effect of *N* and $\rho > 0$ for simulations with $\sigma^2 = 0.1$. Bottom left: Effect of prior $p(\rho|H_1)$ for $\rho \in \{-0.2, -0.1, 0, 0.1, 0.2\}$, $\sigma^2 = 10^{-4}$ and N = 200. Bottom right: Effect of ρ and σ^2 for datasets of size N = 100.

of bivariate normal distribution plus noise (Section III-A); 2) a copula model of dependence (Section III-C); and 3) the abovementioned model of two chaotic systems (Section III-D). In the three models, the intensity of dependence under H_1 was quantified through a parameter (ρ for Sections III-A and III-C, *C* for Section III-D). We varied the parameter and expected $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ to be an increasing function of the intensity of dependence ($|\rho|$ for Section III-A) and Section III-C, *C* for Section III-D.

For all simulations, we focused on $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$, whose value is simple to interpret (a value of *b* means that H_1 is 10^b times more probable than H_0). Its lower and upper bounds are $-\infty$ and $+\infty$, respectively. A summary of the main results can be found in Section III-E.

A. Bivariate Normal Distribution With Noise

1) Model: We considered data generated according to the following model: (X, Y) is a bivariate normal distribution with zero mean and covariance matrix given by

$$\tau^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} = \tau^2 M(\rho) \tag{19}$$

where τ is assumed to be a known parameter. However, we only measured noisy versions (U, V) of (X, Y) related through

$$(U, V) = (X, Y) + (E, F)$$
(20)

with *E* and *F* independent Gaussian variables with zero mean and known variance σ^2 . We observed *N* realizations $(u_n)_{n=1,...,N}$ and $(v_n)_{n=1,...,N}$ of *U* and *V*. To quantify the dependence between *X* and *Y* from these *N* realizations, we used the following two models: 1) H_0 , where *X* and *Y* are uncorrelated ($\rho = 0$) and 3) H_1 where they may be correlated.

The model is analyzed in Section 8 of Supplementary Material, yielding for $p(D|H_0)$

$$(2\pi)^{-\frac{2N}{2}} \left(\sigma^2 + \tau^2\right)^{-\frac{2N}{2}} \exp\left[-\frac{\sum_{n=1}^{N} (u_n^2 + v_n^2)}{2(\sigma^2 + \tau^2)}\right]$$
(21)

and for $p(D|H_1)$

$$(2\pi)^{-\frac{2N}{2}} \int \left| \sigma^2 \mathbf{I} + \tau^2 \mathbf{M}(\rho) \right|^{-\frac{N}{2}} p(\rho|H_1) d\rho$$

$$\times \exp\left[-\frac{1}{2\sigma^2} \operatorname{tr} \left(\mathbf{S} \left\{ \mathbf{I} - \left[\mathbf{I} + \frac{\sigma^2}{\tau^2} \mathbf{M}(\rho)^{-1} \right]^{-1} \right\} \right) \right].$$
(22)

For $p(\rho|H_1)$, we considered a prior that could possibly remove a neighborhood of $\rho = 0$. To this end, we used a general distribution of the form

$$q_{\epsilon}(\rho) = \begin{cases} 0, & \text{for } |\rho| < \epsilon \\ \frac{1}{2(1-\epsilon)}, & \text{otherwise.} \end{cases}$$
(23)

Such a prior imposes $\rho \notin [-\epsilon, \epsilon]$ and is uniform otherwise. For $\epsilon = 0$, this is the usual uniform prior on [-1, 1].

2) Data: We generated data with ρ ranging from -0.9 to +0.9 by increment of 0.1, N ranging from 20 to 200 by increment of 20 as well as 300 and 500, and $\sigma^2 \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$. τ^2 was set to 1. For each particular value of the triplet (ρ, σ^2, N) , we generated M = 1000 samples. For each sample, we computed $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ using priors for ρ of the form given in (23) with $\epsilon \in \{0, 0.01, 0.1, 0.2\}$.

3) *Results:* Results are summarized in Fig. 1. For $\rho = 0$ (corresponding to H_0 true), $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was mostly negative and its value decreased with increasing N and decreasing



Simulation study: functional dependence with noise. Boxplots Fig. 2. (median and [25%, 75%] percentile) of the effect of σ^2 and N on $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ when the true model is either H_1 (top) or H_0 (bottom).

 σ^2 [Fig. 1, top left]. For $\rho \neq 0$ (corresponding to H_1 true), it was mostly positive, and its value increased with increasing N, decreasing σ^2 , and increasing $|\rho|$ [Fig. 1, top and bottom right]. Furthermore, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was observed to behave similarly to a logarithmic function of N for $\rho = 0$ and to a linear function of N for $\rho \neq 0$. Finally, a prior distribution prohibiting small values of $|\rho|$ tended to give more weight to H_0 around $\rho = 0$ [Fig. 1, bottom left]. These results are in line with our predictions regarding the expected behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ when N increases and σ^2 decreases (see beginning of Section III).

B. Functional Dependence With Noise

1) Model: We considered a two-dimensional (2-D) variable (X, Y) where X and Y may be related by a functional relationship. For the sake of simplicity, we considered a linear relationship. More precisely, for dependence (H_1) , we assumed that we had

$$(X, Y) = (T, T) + (E, F)$$
 (24)

with $T \sim \mathcal{N}(0, \tau^2)$. By contrast, for independence (H₀), we assumed

$$(X, Y) = (U, V) + (E, F)$$
(25)

with $U, V \sim \mathcal{N}(0, \tau^2)$. In both cases, E and F are white noise of known variance σ^2 , and τ^2 is assumed to be known. Assume that we observed N realizations $(x_n, y_n)_{n=1,...,N}$ of (X, Y). Note that the intermediary variables (U and V for H_0 ; T for H_1) are not observed. Interestingly, unlike most models of dependence, the description of dependence here requires fewer parameters (the t_n 's) than description of independence (the u_n 's and v_n 's).

We obtain (see Section 9 of Supplementary Material)

$$\ln \frac{p(D|H_1)}{p(D|H_0)} = -\frac{N}{2} \ln(\sigma^2) - \frac{N}{2} \ln(\sigma^2 + 2\tau^2) + N \ln(\sigma^2 + \tau^2) \\ - \sum_{n=1}^{N} \left[\frac{(x_n - y_n)^2}{2\sigma^2(2 + \alpha^2)} + \frac{x_n^2 + y_n^2}{2\tau^2(2 + \alpha^2)} - \frac{x_n^2 + y_n^2}{2\tau^2(1 + \alpha^2)} \right]$$
(26)



Authorized licensed use limited to: IEEE - Staff. Downloaded on April 28,2025 at 11:52:21 UTC from IEEE Xplore. Restrictions apply.

2) Data: We generated data with either model H_0 or model H_1 , N ranging from 20 to 100 by increment of 20. Each time, we generated M = 1000 samples. For each sample, we computed $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ using IFD with (16) and (17).

3) Results: Results are summarized in Fig. 3. As expected, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was found to be 1) negative and a decreasing function of N when H_0 was true, and 2) positive and an increasing function of both $|\rho|$ and N when H_1 was true.

D. Dependence of Two Chaotic Systems

1) Model: To demonstrate the possibility of our measure to quantify the intensity of coupling between two systems, we used the example of two coupled chaotic Rossler oscillators with a small parameter mismatch. Each oscillator $i \in \{1, 2\}$

with $\alpha^2 = \sigma^2 / \tau^2$.



Fig. 3. Simulation study: dependence through copula. Boxplots (median and [25%, 75%] percentile) of the effect of ρ and N on $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ when H_1 is true (top), and of the effect of N when H_0 is true (bottom).

2) Data: We generated data with either model H_0 or model H_1 , N ranging from 20 to 200 by increment of 20 as well as 300 and 500, and $\sigma^2 \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$. τ^2 was set to 1. For each particular value of (H_i, σ^2, N) , we generated M = 1000 samples. For each sample, we computed $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D).$

3) Results: Results are summarized in Fig. 2. When H_0 was true, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was found to be negative, a decreasing function of N, and an increasing function of σ^2 . When H_1 was true, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was found to be positive, an increasing function of N, and a decreasing function of σ^2 . Unlike what was found previously, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ behaved similarly to a linear function of N under both H_0 and H_1 . But these results are again in line with our predictions.

C. Dependence Through Copula

1) Model: We modeled a 2-D variable (X, Y) with marginals equal to gamma distributions, with (α, β) equal to (4, 4) for X and (10, 5) for Y. In H_1 , dependence was modeled through a Student's t copula [23, Sec. 8.3] with 5 degrees of freedom and $\rho \in \{0, 0.2, 0.7\}$. Note that $\rho = 0$ corresponds to uncorrelated, yet dependent variables [26, Sec. 1.16].



Fig. 4. Simulation study: dependence of two chaotic systems. Boxplots (median and [25%, 75%] percentile) of the effect of the coupling parameter *C* when N = 50 (top), and of the effect of *N* with either C = 1, corresponding to H_1 true (middle) or with C = 0, corresponding to H_0 true (bottom).

was characterized by its position (x_i, y_i, z_i) and time derivatives $(\dot{x}_i, \dot{y}_i, \dot{z}_i)$. Coupling was quantified through *C* (*C* = 0 corresponds to no coupling). For more details, see Section 10 of Supplementary Material or [27, Sec. 3.1.4].

2) Data: We simulated data with $C \in \{0, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$. For a given set of parameter values, the trajectory of the system was generated numerically with an explicit Runge–Kutta method and downsampled to one sample per second. Trajectories with $N \in \{10, 20, 50\}$ time points were considered. From each trajectory, M = 100 samples were generated by adding Gaussian white noise with variance $\sigma^2 \in \{10^{-3}, 10^{-2}, 10^{-1}, 1\}$.

3) *Results:* Results are summarized in Fig. 4. As expected, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was globally found to be 1) a decreasing function of N and an increasing function of σ^2 when H_0 was true and 2) an increasing function of both N and C and a decreasing function of σ^2 when H_1 was true. Exceptions to this general trend was the case when C was low and σ^2 large, in which case $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ could first decrease and then increase for increasing N.

E. Summary of Results

In the simulation section, we showed that $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ provided a measure of dependence between *X* and *Y* that had the following properties.

- 1) When H_0 was true, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ typically decreased when the quantity of information available in the data increased (increasing *N*, decreasing σ^2).
- 2) When H_1 was true, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ typically increased when the quantity of information available in the data increased (increasing *N*, decreasing σ^2). Furthermore,

when the intensity of dependence between X and Y was parameterized, $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was found to be an increasing function of this intensity.

IV. REAL-LIFE APPLICATION

Electroencephalography (EEG) is a brain exploration technique that allows to noninvasively record electrical consequences of brain activity. Such recordings are often driven by brain oscillations originating from synchronized neuronal activity. A common procedure for EEG acquisitions is the socalled event-related protocol, where one records how the brain responds (through the evoked response) to a given stimulation over many repetitions, called trials. For some types of protocols, the stimulus may consistently induce synchronization of brain activity, which translates into a phenomenon called phase resetting. In this case, the phase of the signal (quantified, e.g., through time-frequency analysis) in a certain time window after the stimulus remains consistent over trials. It θ_n is the phase quantified for trial n, n = 1, ..., N, phase consistency has typically been quantified using intertrial phase coherence (ITC) [28], [29], which, in circular statistics, is the mean resultant length \overline{R} of the sample (θ_n) [30, Sec. 2.3.1]

$$\overline{R} = \left| \frac{1}{N} \sum_{n=1}^{N} e^{i\theta_n} \right|.$$

In our framework, we can propose an alternative measure of the dependence between the stimulus and the brain. More specifically, we consider two competing models H_0 and H_1 , where H_0 assume that θ_n is uniformly distributed on the circle, while H_1 assumes that θ_n has a von Mises distribution with mean direction μ and concentration parameter κ [30, Sec. 3.5.4]. Using standard prior distributions for μ and κ , we obtain the following measure of dependence (see Section 11 of Supplementary Material)

$$\mathfrak{B}_{\log r}(\mathcal{X}, \mathcal{Y}|D) = \log_{10} \frac{p(H_1)}{p(H_0)} + \log_{10} \left[\int \frac{\kappa I_0(N\overline{R}\kappa)}{\left(1 + \kappa^2\right)^{\frac{3}{2}} I_0(\kappa)^N} \, d\kappa \right].$$

In the following, we assume $p(H_1) = p(H_0) = 1/2$. The integral can be computed numerically for any value of N and \overline{R} .

Results are summarized in Fig. 5. $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ was found to be an increasing function of \overline{R} , a decreasing function of Nfor lower values of \overline{R} ($\overline{R} \leq 0.1$), and an increasing function of N for larger values of \overline{R} ($\overline{R} \geq 0.3$). For intermediate values of \overline{R} , $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ first decreased, then increased. To further investigate this change in monotonicity, we computed for different values of \overline{R}

$$N_0(R) = \operatorname{argmin}_N \mathfrak{B}_{\operatorname{logr}}(\mathcal{X}, \mathcal{Y}|D).$$

For fixed \overline{R} , $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ decreased for $N \leq N_0(\overline{R})$ and increased for $N \geq N_0(\overline{R})$. Since $2N\overline{R}^2$ is approximately χ_2^2 for large N under the assumption of uniform phase (i.e., H_0) [30, Sec. 4.4], we have $E(\overline{R}^2|H_0) \approx 1/N$, that is, we can



Fig. 5. Real-life application. $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$ as a function of \overline{R} for various values of N (top) and as a function of N for various values of \overline{R} (middle). Bottom panel: contour plot of $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D)$, together with $[N_0(\overline{R}), \overline{R}]$ (black solid line and circles) and $(N, 1/\sqrt{N})$ (black dashed line).

expect \overline{R} to have values of the order of $1/\sqrt{N}$ (still under H_0). We empirically noticed that $N_0(\overline{R})$ was smaller than, yet close to $1/\sqrt{N}$.

V. GENERAL PROPERTIES OF $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$

Various researchers have tried to define a set of properties that a good measure of dependence should respect [4], [8], [31], [32], [33], [34], [35]. The properties of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ exhibited so far either in the calculations (Section II) or the stimulation study (Section III) are closely related to these properties.

- 1) It is well defined for any pair of (either continuous or discrete) variables X and Y as long as the corresponding models H_0 and H_1 are.
- 2) It is symmetrical in X and Y as long as both H_0 and H_1 share this property.
- 3) It reaches its minimum value when X and Y are known to be independent as described by H_0 .
- 4) It reaches its maximum value when X and Y are known to be dependent as described by H_1 .

- 5) It is an increasing function of mutual information, when such a measure makes sense.
- 6) It is an increasing function of the linear correlation coefficient in the case where (X, Y) follows a bivariate normal distribution.

Properties 3 and 4 are quite restrictive, since we assume that we *know* what the underlying model is. Indeed, we usually do not know for sure whether H_0 or H_1 is true. In this case, we must set $Pr(H_0) \neq 0$ and $Pr(H_1) \neq 0$. Results from both the theoretical calculations and the simulation studies indicate that the following original properties also hold asymptotically.

- 3') If H_0 is true, then $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is a decreasing function of N which tends to its lower bound as $N \to \infty$.
- 4') If H_1 is true, then $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is an increasing function of N which tends to its upper bound as $N \to \infty$.
- In case of model misspecification,
 ^B(X, Y|D) behaves
 as if the true model were the one closer to the true
 generative model (in the sense of Kullback–Leibler
 divergence).

Note that, from all the results presented in this manuscript, Property 3' seemed to hold even for small values of N, while Property 4' sometimes required larger values of N to hold and followed an initial stage where $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ decreased. Finally, we observed the following empirical properties from the simulation studies.

- The effect of noise was the reverse of the effect of N: it tended to increase B(X, Y|D) when H₀ was true and to decrease it when H₁ was true.
- 9) When H₁ was true and the intensity of dependence was quantified by a parameter, B(X, Y|D) was an increasing function of this intensity.

To our knowledge, it is the first time that these properties are considered as potentially desirable features for a measure of dependence. This is further discussed in the next section.

VI. DISCUSSION

Summary: Quantification of dependence between two systems is still an open issue in the general case. We here proposed a general measure of dependence $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ between two systems \mathcal{X} and \mathcal{Y} for a given dataset D based on the Bayesian comparison of two models, one of independence H_0 and another one of dependence H_1 . Dependence within D was then quantified as the posterior probability of H_1 given D, $\mathfrak{B}_{\mathrm{pr}}(\mathcal{X}, \mathcal{Y}|D) = \mathrm{Pr}(H_1|D)$, or a strictly increasing function of it. We calculated the value of $\mathfrak{B}_{lnr}(\mathcal{X}, \mathcal{Y}|D) = \ln \mathfrak{B}_{pr}(\mathcal{X}, \mathcal{Y}|D)$ in particular cases: when the model distribution is known, when it belongs to a known parametric family with unknown parameters, when H_0 is nested in H_1 (including common models, such as maximum-entropy distributions, multivariate normal distributions, and bivariate discrete distributions), and when dependence is modeled through a copula. Using simulations, we investigated the behavior of $\mathfrak{B}_{logr}(\mathcal{X}, \mathcal{Y}|D) =$ $\log_{10} \mathfrak{B}_{pr}(\mathcal{X}, \mathcal{Y}|D)$ in the presence of noise and when H_1 had a parameter coding for the intensity of dependence. We also provided an application of our framework to neuroscience and neuroimaging. Finally, we stated some key general properties of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$. While some of these properties are typical of what is expected from a valid measure of dependence, others are novel and naturally appear as desirable features for $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$.

Existing Work: The theoretical results introduced here relating the posterior distribution and its variants to mutual information (Section II) bring together two existing lines of research from standard statistical theory and information theory. On the one hand, the behavior of the model marginal likelihood for large N has been studied in depth and the theoretical underpinnings of such calculations go back to the development of the BIC and the minimum description length (MDL). On the other hand, the asymptotic bias of empirical mutual information estimators is also well known. The connection between Bayesian posterior distribution and mutual information was first noticed in particular cases by [15], [16], and [17]. We here showed that this connection actually holds in the more general setting of nested models (Section II-E). While these results involve developments that are rather standard in statistical theory and information theory, they shed some important light on the relevance of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ as a valid measure of dependence.

To our knowledge, no general theory exists regarding Bayesian model comparison for non-nested models. Our results from copula modeling of dependence (Sections II-F and III-C) suggested that similar results regarding the behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ might hold in that particular case.

Difference of Behavior Under H_0 and H_1 : It has to be underlined that the roles of H_0 (independence) and H_1 (dependence) are not symmetrical. Often, H_0 is nested in H_1 or, equivalently, H_1 includes H_0 as a particular case (parameters set to particular values, e.g., usually 0). This has the following consequence for Bayesian model comparison. In the case of independence, a H_1 with likely parameter values becoming increasingly closer to, e.g., 0 cannot be fully ruled out. By contrast, in the case of dependence, H_0 becomes increasingly unlikely as N increases. This translates into a typically different behavior of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$: a logarithmic decrease under H_0 but a linear increase under H_1 . Such a behavior, which was observed in all our computations as well as our first example in the simulation study, was associated with a connection with mutual information. By contrast, the second example of our simulation study (Section III-B) exhibited linear behaviors both when H_0 and H_1 were true, together with a connection with mutual information that remains to be determined.

Another difference between H_0 and H_1 can be seen in the monotonic property of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ with respect to N. As detailed in various places in the manuscript, $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ under H_0 was found to be a decreasing function of N even for moderate to small values of N. By contrast, under H_1 , it cannot be ruled that $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ has an initial stage where it is actually a decreasing function of N, before becoming an increasing function. We interpret it as a consequence of the fact that Bayesian analysis tries to find a compromise between the strength of dependence and model parsimony, in line with the usual interpretation of Bayesian inference as providing a quantitative implementation of Occam's razor. How this nonmonotonic behavior fits in our general framework remains to be further clarified.

Measures of Dependence: The main point of the present work is that $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is a valid measure of dependence. Quantification of dependence is still a field of ongoing research, whose objective is to provide a measure that quantifies the departure of $f_{XY}(x, y)$ from independence, i.e., from $f_X(x)f_Y(y)$. To our knowledge, the present work is the first one to advocate that the posterior probability in a specifically designed Bayesian model comparison analysis can be considered as a valid measure of dependence. Importantly, $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$, as a result of a Bayesian analysis, quantifies by construction the evidence for dependence between \mathcal{X} and \mathcal{Y} as modeled by H_1 (compared to H_0) and observed in D. It is what we would call an *inferential* measure of dependence, in that it is both model-based, as it incorporates information from (probabilistic) models H_0 and H_1 , and *data-driven*, as its value reflects the content of a dataset D. As a consequence, the same value of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ can be obtained for two very different scenarios: 1) either a large amount of data about weakly dependent variables or 2) a small amount of data about strongly associated variables. Yet, for a given system, changes in $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ can be interpreted unambiguously: when the information content of D increases (e.g., with increasing size or decreasing noise), the measure becomes increasingly closer to the boundary of its definition domain that corresponds to the correct ideal situations, in agreement with Properties 3, 3', 4, and 4' of Section V. Importantly, this is not a direct consequence of the fact that we used a Bayesian model comparison analysis but is the result of a selective choice of the measure. For instance, a nonmonotonic function of $Pr(H_1|D)$, such as $(1/N) \ln Pr(H_1|D)$, which behaves in a fashion very similar to mutual information, would not qualify as an inferential measure of dependence (see Section-12 of Supplementary Material). We will come back to this point below.

In parallel to the development of new measures of dependence, there has also been active research regarding the properties that a good measure of dependence should respect [4], [8], [31], [32], [33], [34], [35]. In Section V, we provided properties that hold for $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$. Properties 1–6 are based on previous descriptions of how ideal measures of dependence should behave, while Properties 3', 4', 7, 8, and 9 have been introduced in this manuscript and are specific to the expected behavior of what we coined inferential measures of dependence.

Some authors also consider that a good measure of dependence should be a metric [33], [34], [35]. We did not check this requirement, and it would be of interest to see whether it can be met by $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$. Still, according to the main properties of a dependence measure, the more two variables are dependent, the larger the measure. In classification terminology, this makes a dependence measure closer to a similarity measure than to a distance measure. To our knowledge, this is in agreement with only one instance of measure proposed in [36], which is indeed a decreasing function of dependence and, therefore, does not meet Properties 3–6.

Posterior Probability and Mutual Information: In this manuscript, we made several connections between $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ and the plug-in estimator of mutual information $\hat{I}(X, Y)$,

showing that $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is often, even approximately, an affine function of \hat{I} . Still, there are some major differences between $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ and $\hat{I}(X, Y)$, and more generally, mutual information, that need to be stated and clarified.

First, we argue that $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is more general than mutual information in two senses. First, while there are many cases where there is a (direct or indirect) connection between $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ and mutual information, there are also cases where such a connection may not make sense or does not exist (see, e.g., the simulation study, Section III and Section 13 of Supplementary Material). In such cases, we advocate that the use of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ —instead of mutual information—is justified by the framework we introduced in the present manuscript and still makes sense. Second, unlike mutual information, whose value taken by an estimator on a given dataset of finite size does not have any meaning, the values taken by Bayesian measures of dependence have a simple interpretation, as they quantify the evidence in favor of dependence for the dataset under consideration.

Another major difference between $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ and $\hat{I}(X, Y)$ is that they are measures of dependence of different nature. $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is obtained through hypothesis testing and quantifies the evidence of H_1 against H_0 . By contrast, mutual information quantifies the theoretical level of dependence contained in a model H_1 compared to H_0 . It is usually a function of model parameters that need to be estimated to provide an estimate of mutual information. As a consequence, when H_1 is true (dependence), 1) E[I(X, Y)] is not necessarily an increasing function of N as $N \rightarrow \infty$ and 2) it does not tend to the upper bound of its range $(+\infty)$ as $N \to \infty$. More precisely, in the case of independence, H_0 can often be associated with a zero mutual information (Sections II-C-II-E). In this case, any valid estimator of mutual information will tend to 0 as $N \rightarrow \infty$. Since 0 is the lower bound of mutual information, Properties 3 and 3' can be expected to roughly hold. By contrast, if H_1 is true (dependence), any valid estimator of mutual information will tend to the theoretical value of mutual information (which is in general strictly lower than its upper bound, $+\infty$), and changes in the estimator values will be mainly due to statistical fluctuations around this theoretical value. As a consequence, Properties 4 and 4' of Section V are not respected.

A model-centered measure aims at quantifying the theoretical level of dependence between two variables entailed by, or contained in, a model (and not data). It ranks models, from one(s) with the least dependence (usually independence, for which the measure reaches its lower bound) to one(s) with the most dependence (for which the measure reaches its upper bound). Data are then used to infer this theoretical level of dependence *from* them. Such a feature is not specific to mutual information, and the family of model-centered measures include many other existing measures of dependence, including all those mentioned in the introduction (maximal correlation coefficient [3], [4], the mixed derivative measure of marginal interaction [1, Sec. 2.3], Hoeffding's procedure [5], distance correlation [6], circular correlation [8], and Hilbert– Schmidt information criterion [9]). *Limitations of Method:* The current approach strongly relies on Bayesian model comparison. As a consequence, it has limitations that are typical of that kind of approaches and are related to the choice of the model, the choice of the priors, and the computation of the marginal model likelihoods.

First, our approach requires the specification of two models, one for independence (H_0) and one for dependence (H_1) . Of course, it cannot be ruled out that either, or both models are incorrect, hence the importance of considering model misspecification. In that case, it is possible to show (see Section II-G) that, under specific assumptions, $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ will behave as if the true model were the one that is closer (in the sense of Kullback-Leibler divergence) to the true generative one. This is in line with existing general results regarding Bayesian model comparison [19, Sec. 7.27]. Note that the choice of H_0 is often dictated by the type of data considered. It is somewhat made easier by the (strong) constrain of (1). By contrast, the choice for H_1 should take into account both the type of data and the potential structure of dependence. As a consequence, we expect this choice to be more complex and prone to successive adjustments.

Also, Bayesian analyses require the introduction of prior distributions. Here, we needed prior information regarding both the relative plausibility of H_0 and H_1 as well as the potential values of the parameters for both models. As data size grows, we expect the respective priors for H_0 and H_1 to have vanishing impact on $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$, unless there is prior evidence that one model is overwhelmingly more plausible than the other. By contrast, choosing the priors on the model parameters is more problematic. Noninformative priors, which are commonly used in Bayesian parameter estimation for the sake of simplicity, are strongly advised against for model comparison. Conjugate priors, which are often used, might not correctly represent the prior information at hand, while using tailored priors might lead to an intractable $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$. See [15] for a discussion of the choice of the prior on the covariance matrix for the multivariate normal model. Unlike Bayesian parameter estimation, where the impact of the prior vanishes for large N, Bayesian model comparison is influenced by the choice of the model parameter prior for any data size. Still, the manuscript provides two results that specify the role of this type of prior. First, we theoretically showed that, with increasing data size, it is the dimension of the parameter space that matters rather than the distribution of its values in that parameter space (see Sections II-D and II-E), in line with the usual Laplace and BIC approximations. Since Bayesian analysis can involve very complex models with many parameters in high dimension, in particular when modeling dependence in H_1 , these asymptotic results may be of limited relevance in some applications. In these cases, one needs to go back to the exact formulas with finite N. The parameter prior may then have an effect on $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$, an effect that is important to quantify, e.g., through sensitivity analysis using different prior distributions. For instance, we empirically showed on simulated data [Section III-A, see in particular $q_{\epsilon}(\rho)$ and (23)] that having a more specific prior had limited influence when H_1 was true and the model parameter

was indeed in the correct range, while helping ruling out H_1 when H_0 was true.

Finally, the broad applicability of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ will partly depend on the ease with which ratios of marginal model likelihoods can be computed. Exact calculation of marginal model likelihoods is usually very difficult, from both an analytical and a computational point of view. Analytical calculations often rely on simple models combined with conjugate priors (as in the present manuscript), tricks, such as the Savage-Dickey density ratio [37], or approximations, e.g., Laplace approximation [20], [21] or the BIC [22]. Exact computational approaches are based on numerical integration, which is time-intensive and suffers from the curse of dimensionality. A wealth of approximate methods have been devised, with specific advantages and limits, which make them more or less suitable depending on the context: the acceptance ratio method and thermodynamic integration [38]; importance sampling, bridge sampling, and path sampling [39, Sec. 5]; reversible jump Markov chain Monte Carlo (RJ-MCMC) [40]; and nested sampling [41].

Generalization to Several Variables: In the present manuscript, we focused on the quantification of dependence between two (potentially multidimensional) variables. The framework can very easily be expanded to take into account dependence between several variables. The general framework is the same, and was detailed in [16]. The connection between Bayesian measures of dependence and mutual information still holds, with mutual information replaced by a generalization to several variables known as total correlation [42], multivariate constraint [43], δ [44], or multiinformation [45].

Questioning the Notion of Dependence: Importantly, the use of our framework made it clear that independence itself is not enough to provide an unambiguous measure of dependence. Indeed, there was a need to introduce both an alternative model H_1 , in which a potential structure of dependence was introduced, and data, which where used to quantify dependence. Both the model of dependence and the data have a key influence on the resulting dependence measure.

It cannot be excluded that data that appear to be associated to independent variables are actually dependent but with a dependence structure that is not the one introduced in H_1 . This could appear as a weakness of the method, but we believe it is rather a strength of it: while the underlying assumptions (in particular regarding the model of dependence H_1) are implicit in many methods, here they have to be clearly stated and translated into operational form to be able to perform a Bayesian analysis. We argue that this makes the model easy to falsify and, as a consequence, to improve. It is not uncommon that variables that were believed to be independent are later found dependent based on a different model of the dependence pattern. For instance, in neuroimaging, where the authors have extended experience, there has been a major interest in the subfield of functional connectivity analysis, where one tries to use brain imaging data to quantify the dependence between brain regions. Unsurprisingly, functional connectivity is quantified differently based on the imaging modality. In functional magnetic resonance imaging (fMRI) [46], [47], the data are commonly assumed to follow

a normal distribution and dependence is often quantified through pairwise correlation [48], [49]. By contrast, in EEG and magnetoenccephalography (MEG) [50], pairwise correlation is considered as a poor measure of dependence, and other, more adapted measures of functional connectivity have been proposed, based on models of circular or oscillating data [51], [52]. Another example is the real-life application of Section IV), where the two advantages of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ appear clearly: 1) the underlying assumptions are made apparent (uniform versus von Mises phase) and 2) these assumptions can be checked, falsified, and possibly changed for more realistic modeling.

Paralleling what was said for H_1 , it cannot be excluded that variables that are associated with low or decreasing values of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ based on small datasets would actually yield large or increasing values of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ with larger data sets. As mentioned above, Bayesian analysis comes with a built-in quantitative implementation of Occam's razor which performs a compromise between the strength of dependence and model parsimony. As a consequence, data originating from weakly dependent models may at first yield low or decreasing values of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$. Such a behavior emphasizes the importance of considering $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ as a measure of the *evidence for dependence in a given dataset with a given model of dependence.*

VII. CONCLUSION

The present work is the first one to advocate that the posterior probability (or any strictly increasing measure thereof) resulting from a specifically designed Bayesian model comparison analysis can be considered as a valid measure of dependence. We showed that such a framework provided a family of measures, denoted $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ that quantify the information contained in *D* in favor of H_1 versus H_0 . As such, they quantify the evidence for—or credibility of—dependence between \mathcal{X} and \mathcal{Y} as modeled by H_1 (compared to H_0) and observed in *D*. All measures in this family shared the following key asymptotic properties in a wide range of situations.

- 1) Under H_0 , $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ is a decreasing function of N which tends to its lower bound when $N \to \infty$;
- Under H₁, 𝔅(𝔅, 𝔅|D) is an increasing function of N for N large enough (after a potential initial stage of decrease) and tends to its upper bound when N → ∞;
- B(X, Y|D) is an increasing function of some common existing measures of dependence, such as correlation, mutual information, the minimum discrimination information, and the log-likelihood ratio criterion for testing independence;
- In case of model misspecification,
 ^B(X, Y|D) behaves
 as if the true model were the one closer to the true
 generative model.

Empirically, we also showed that 1) increasing noise had the opposite effect on $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ to increasing *N*, in that it drew the measure away from the expected (lower for H_0 , upper for H_1) bound and 2) when H_1 had a parameter coding for the intensity of dependence, $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ was an increasing function of this intensity. Finally, the value computed for a

given dataset of finite size has a direct, exact interpretation as (a strictly increasing function of) the probability that \mathcal{X} and \mathcal{Y} be dependent (as described by H_1) for that given dataset. Our objective is now to show the generality and versatility of $\mathfrak{B}(\mathcal{X}, \mathcal{Y}|D)$ as an inferential measure of dependence.

ACKNOWLEDGMENT

The authors are grateful to the reviewers for their comments which substantially improved the quality of the manuscript.

REFERENCES

- J. Whittaker, Graphical Models in Applied Multivariate Statistics. Chichester, U.K.: Wiley, 1990.
- [2] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Telecommunications and Signal Processing). Hoboken, NJ, USA: Wiley, 1991.
- [3] H. Gebelein, "Das statistische problem der korrelation als variationsund eigenwertproblem und sein zusammenhang mit der ausgleichsrechnung," Z. Angew. Math. Mech., vol. 21, no. 6, pp. 364–379, 1941.
- [4] A. Rényi, "On measures of dependence," Acta Math. Acad. Sci. Hung., vol. 10, nos. 3–4, pp. 441–451, 1959.
- [5] W. Hoeffding, "A non-parametric test of independence," Ann. Math. Statist., vol. 19, no. 4, pp. 546–557, 1948.
- [6] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, "Measuring and testing dependence by correlation of distances," *Ann. Statist.*, vol. 35, no. 6, pp. 2769–2794, 2007.
- [7] D. Edelmann, T. F. Móri, and G. J. Székely, "On the relationships between the Pearson and the distance correlation coefficients," *Statist. Probab. Lett.*, vol. 169, Feb. 2021, Art. no. 108960.
- [8] P. E. Jupp and K. V. Mardia, "A general correlation coefficient for directional data and related regression problems," *Biometrika*, vol. 67, no. 1, pp. 163–173, 1980.
- [9] A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf, "Measuring statistical dependence with Hilbert-Schmidt norms," in *Proc. Int. Conf. Algorithmic Learn. Theory*, 2005, pp. 63–78.
- [10] T. W. Anderson, An Introduction to Multivariate Statistical Analysis (Wiley Publications in Statistics). Hoboken, NJ, USA: Wiley, 1958.
- [11] S. Kullback, *Information Theory and Statistics*. Mineola, NY, USA: Dover, 1968.
- [12] L. Paninski, "Estimation of entropy and mutual information," *Neural Comput.*, vol. 15, no. 6, pp. 1191–1253, 2003.
- [13] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating mutual information," *Phys. Rev. E*, vol. 69, Jun. 2004, Art. no. 66138.
- [14] A. Kraskov, H. Stögbauer, and P. Grassberger, "Erratum: Estimating mutual information [Phys. Rev. E 69, 066138 (2004)]," Phys. Rev. E, vol. 83, Jan. 2011, Art. no. 19903.
- [15] G. Marrelec, A. Messé, and P. Bellec, "A Bayesian alternative to mutual information for the hierarchical clustering of dependent random variables," *PLoS One*, vol. 10, no. 9, 2015, Art. no. e0137278.
- [16] G. Marrelec and A. Giron, "Automated extraction of mutual independence patterns using Bayesian comparison of partition models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 7, pp. 2299–2313, Jul. 2021.
- [17] D. R. Wolf, "Mutual information as a Bayesian measure of independence," 1994, arXiv:comp-gas/9511002.
- [18] I. J. Good, "Studies in the history of probability and statistics. XXXVII A. M. Turing's statistical work in World War II," *Biometrika*, vol. 66, no. 2, pp. 393–296, 1979.
- [19] A. O'Hagan and J. Forster, Kendall's Advanced Theory of Statistics: Vol. 2B: Bayesian Inference. London, U.K.: Arnold, 2004.
- [20] L. Tierney and J. B. Kadane, "Accurate approximations for posterior moments and marginal densities," J. Amer. Statist. Assoc., vol. 81, no. 393, pp. 82–86, 1986.
- [21] A. E. Gelfand and D. K. Dey, "Bayesian model choice: Asymptotics and exact calculations," *J. Roy. Statist. Soc. Ser. B*, vol. 56, no. 3, pp. 501–514, 1994.
- [22] G. Schwarz, "Estimating the dimension of a model," Ann. Statist., vol. 6, no. 2, pp. 461–464, 1978.
- [23] D. Ruppert and D. S. Matteson, *Statistical and Data Analysis for Financial Engineering* (Springer Texts in Statistics). New York, NY, USA: Springer, 2015.

- [24] H. Joe, Multivariate Models and Multivariate Dependence Concepts (CRC Monographs on Statistics & Applied Probability). Boca Raton, FL, USA: Chapman Hall, 1997.
- [25] J. Ma and Z. Sun, "Mutual information is copula entropy," *Tsinghua Sci. Technol.*, vol. 16, no. 1, pp. 51–54, 2011.
- [26] S. Kotz and S. Nadarajah, *Multivariate t Distributions and their Applications*. Cambridge, U,K: Cambridge Univ. 2004.
- [27] S. Boccaletti, J. Kurths, G. Osipov, D. L. Valladares, and C. S. Zhou, "The synchronization of chaotic systems," *Phys. Rep.*, vol. 366, pp. 1–101, Aug. 2002.
- [28] S. Makeig, S. Debener, J. Onton, and A. Delorme, "Mining event-related brain dynamics," *Trends Cogn. Sci.*, vol. 8, no. 5, pp. 204–210, 2004.
- [29] J. Benhamou, M. Le Van Quyen, and G. Marrelec, "Time-frequency analysis of event-related brain recordings: Connecting power of evoked potential and inter-trial coherence," *IEEE Trans. Biomed. Eng.*, vol. 70, no. 5, pp. 1599–1610, May 2023.
- [30] K. V. Mardia and P. E. Jupp, *Directional Statistics* (Wiley Series in Probability and Statistics). Chichester, U.K.: Wiley, 2000.
- [31] L. A. Goodman and W. H. Kruskal, "Measures of association for cross classifications," *J. Amer. Statist. Assoc.*, vol. 49, no. 268, pp. 732–764, 1954.
- [32] W. J. Hall, "On characterizing dependence in joint distributions," in *Essays in Probability and Statistics*. Chapel Hill, NC, USA: Univ. North Carolina, 1967.
- [33] E. Maasoumi and J. Racine, "Entropy and predictability of stock market return," J. Econom., vol. 107, pp. 1–2, Mar. 2002.
- [34] C. W. Granger, E. Maasoumi, and J. Racine, "A dependence metric for possibly nonlinear processes," *J. Time Series Anal.*, vol. 25, no. 5, pp. 649–669, 2004.
- [35] A. Dionisio, R. Menezes, and D. A. Mendes, "Entropy-based independence test," *Nonlinear Dyn.*, vol. 44, pp. 351–357, Jun. 2006.
- [36] A. Kraskov, H. Stögbauer, R. G. Andrzejak, and P. Grassberger, "Hierarchical clustering using mutual information," *Europhys. Lett.*, vol. 70, no. 2, pp. 278–284, 2005.
- [37] J. M. Dickey, "The weighted likelihood ratio, linear hypotheses on normal location parameters," *Ann. Math. Statist.*, vol. 42, no. 1, pp. 204–223, 1971.
- [38] R. M. Neal, "Probabilistic inference using Markov Chain Monte Carlo methods," Dept. Comput. Sci., Univ. Toronto, Rep. CRG-TR-93-1, 1993.
- [39] M.-H. Chen, Q.-M. Shao, and J. G. Ibrahim, *Monte Carlo Methods in Bayesian Computation*. New York, NY, USA: Springer, 2000.
- [40] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.
- [41] J. Skilling, "Nested sampling," in Proc. Conf. Bayesian Inference Maximum Entropy Methods Sci. Eng., 2004, pp. 395–405.
- [42] S. Watanabe, "Information theoretical analysis of multivariate correlation," *IBM J. Res. Develop.*, vol. 4, no. 1, pp. 66–82, 1960.
- [43] W. R. Garner, Uncertainty and Structure as Psychological Concepts. New York, NY, USA: Wiley, 1962.
- [44] H. Joe, "Estimation of entropy and other functionals of a multivariate density," Ann. Inst. Statist. Math., vol. 41, no. 4, pp. 683–697, 1989.
- [45] M. Studeny, "Complexity of structural models," in Proc. Joint Session 6th Prague Conf. Asymptot. Stat. 13th Prague Conf. Inf. Theory, Stat. Decision Funct. Random Processes, 1998, pp. 23–28.
- [46] S. A. Huettel, A. W. Song, and G. McCarthy, Functional Magnetic Resonance Imaging. Sunderland, U.K.: Sinauer, 2004.
- [47] R. B. Buxton, Introduction to Functional Magnetic Resonance Imaging.Cambridge, U.K.: Cambridge Univ., 2009.
- [48] G. Marrelec, P. Bellec, and H. Benali, "Exploring large-scale brain networks," J. Physiol. Paris, vol. 100, pp. 171–181, Oct. 2006.
- [49] P. Bellec, A. Messé, D. Coynel, V. Perlbarg, H. Benali, and G. Marrelec, "Resting-state brain networks in functional MRI," in *Brain Mapping: From Neural Basis of Cognition to Surgical Applications*, H. Duffau, Ed., Vienna, Austria: Springer, 2012, pp. 361–372.
- [50] R. Hari and A. Puce, *MEG-EEG Primer*. Oxford, U.K.: Oxford Univ., 2017.
- [51] G. Marrelec, J. Daunizeau, M. Pélégrini-Issac, J. Doyon, and H. Benali, "Conditional correlation as a measure of mediated interactivity in fMRI and MEG/EEG," *IEEE Trans. Signal Process.*, vol. 53, pp. 3503–3516, Sep. 2005.
- [52] E. van Diessen et al., "Opportunities and methodological challenges in EEG and MEG resting state functional brain network research," *Clin. Neurophysiol.*, vol. 126, no. 8, pp. 1468–1481, 2015.

Guillaume Marrelec received the M.S. degree in engineering and applied mathematics jointly from the École Centrale Paris, Gif-sur-Yvette, France, and the Universität Stuttgart, Stuttgart, Germany, in 1999, and the Ph.D. degree in medical imaging from the Université Paris XI, Orsay, France, in 2003.

He is currently a Senior Research Scientist with the Institut national de la santé et de la recherche médicale, Paris, France, and works with the Laboratoire d'imagerie biomédicale, Sorbonne Universite, Paris. His research interests include brain imaging, functional connectivity, information theory, Bayesian analysis, and interactions. Alain Giron received the Ph.D. degree in biomedical engineering from the Universite de technologie de Compiègne (UTC), Compiègne, France, in 1990.

He is currently a Senior Research Engineer with the Institut national de la sante et de la recherche medicale, Paris, France, and works with the Laboratoire d'imagerie biomédicale, Sorbonne Universite, Paris. His research interests include medical imaging, and statistics and artificial intelligence.