Bootstrap based test for the unimodality of estimated Hurst exponents. Performance assessment in a high-dimensional analysis setting

Charles-Gérard LUCAS[1] Patrice ABRY[1] Herwig WENDT[2] Gustavo DIDIER[3] Oliver OREJOLA[3]

[1]ENSL, CNRS, Laboratoire de physique, F-69342 Lyon, France.

[2]IRIT, Univ. Toulouse, CNRS, Toulouse, France.

[3]Math. Dept., Tulane University, New Orleans, USA.

SUPPORTED BY PHD GRANT DGA/AID (NO 01D20019023), ANR-18-CE45-0007 MUTATION.

Résumé – Dans les applications du monde moderne, les systèmes de grande taille sont en général monitorés par de nombreux capteurs, impliquant l'analyse d'un grand nombre de signaux, souvent caractérisés par des dynamiques temporelles invariantes d'échelle. Ce travail vise à construire et valider une procédure bootstrap permettant, à partir d'une seule observation de taille finie, de tester l'unimodalité de la distribution associée des exposants de Hurst estimés. L'analyse de ces performances est conduite dans un cadre d'asymptotique de grandes dimensions (nombre de composantes et taille d'échantillons croissent conjointement).

Abstract – In modern real-world applications, large systems are in general monitored by a large number of sensors, hence entailing the joint analysis of numerous time series, often showing scale-free temporal dynamics. The present work aims to construct and assess a bootstrap based procedure permitting to test, from a single finite size observation, the unimodality of the corresponding distibution of estimated Hurst exponents. Performance analysis is conducted in high-dimensional asymptotic settings, where the number of components and sample size increase jointly.

1 Introduction

Context. Scale-free dynamics are involved in a wide range of applications. They imply that joint temporal dynamics of several time series cannot be reduced to characteristic scales. Instead, scale-free analysis focus on dependencies within a wide range of scales and the key issue lies in the identification of scaling exponents that govern the temporal dynamics within that range of scales. Up to now, scale-free dynamics has mostly been studied in a univariate setting whereas in modern applications, one same system can be monitored by many sensors resulting in a large number of time series that need to be analyzed jointly, such as in neurosciences or climate studies. To overcome this problem, multivariate self-similarity analysis provides a collection of Hurst exponents that jointly characterize multivariate scale-free time series. However, the potentially large number of Hurst exponents raises the crucial issue of testing unimodality vs. multimodality in estimated Hurst exponents, i.e., the presence of one or several different Hurst exponents. Furthermore, to be realistic with respect to applications, the performance of such tests needs to be assessed in high-dimensional settings, where the number of time series is not fixed but grows with sample size.

Related works. Fractional Brownian motion (fBm), the most common model for self-similarity, has recently been extended to the multivariate setting by the so-called operator fractional Brownian motion (ofBm) [2, 5]. Based on this model, a robust eigen-wavelet estimation procedure for the Hurst exponent vector provides a collection of as many estimates as the observed time series constituting the data [1, 4, 9]. In practice, it is often important to estimate whether one single or several different Hurst exponents are actually driving the joint dynamics of the observations. A solution for testing Hurst unimodality has been proposed in [9] and its performance studied in low-dimensional settings where the number of components is fixed, but sample-size increases. The proposed strategy combines eigen-wavelet analysis with a block-bootstrap resampling. An alternative test was proposed in [11], exploiting the properties of large random covariance matrix eigenvalue distribution [3]. Its performance were assessed in a large dimensional settings. Yet the weakness of that proposed procedure stemmed from the fact that the properties of the corresponding test statistics were estimated from multiple observations, thus excluding the use of the test without prior benchmarking relying on simulated ofbm.

Goals, contributions and outline. The present work aims to design a bootstrap procedure to test Hurst unimodality from a single observation of multivariate data adapted to the context of high-dimensional asymptotic limits (number of components, sample size and analysis scales go to infinity jointly). To that end, Section 2 introduces a multivariate self-similarity model (as a particular case of ofBm) and the corresponding multivariate eigen-wavelet estimation procedure for the Hurst exponent vector. Section 3 defines the high-dimensional asymptotic limits and recalls the high-dimensional asymptotic behavior of the distribution of the Hurst exponent vector estimator. The core contribution of this work, described in Section 4, is to contruct a test procedure for Hurst unimodality based on a wavelet-domain block-bootstrap scheme. Section 5 reports test performance assessment, in high dimensional settings, from Monte Carlo experiments conducted on synthetic finite-size ofBm. Reported simulations *(i)* explore numerically conditions where high-dimensional asymptotic limits can be exploited, *(ii)* show that the bootstrap procedure accurately reproduces the null hypothesis of the test and *(iii)* quantify the power of the test.

2 Multivariate self-similarity

2.1 Multivariate fractional Brownian motion

In the present work, we consider the M-fBm, a particular case of *operator fractional Brownian motion* (ofBm), a reference model for multivariate self-similarity [5]. It consists of a collection of M fBm $X \triangleq \{X_1(t), \ldots, X_M(t)\}_{t\in\mathbb{R}}$ of covariance Σ_X, with possibly different Hurst exponents $\underline{H} = (H_1, \ldots, H_M)$, mixed by a $M \times M$ real-valued invertible matrix W,

$$Y \triangleq \{Y_1^{\underline{H},\Sigma_X,W}(t), \ldots, Y_M^{\underline{H},\Sigma_X,W}(t)\}_{t\in\mathbb{R}} \triangleq WX. \quad (1)$$

2.2 Multivariate wavelet eigenvalue regression

Estimating the Hurst exponent vector $\underline{H} = (H_1, \ldots, H_M)$ from an observed M-variate time series Y, the core of multivariate self-similarity analysis, has been addressed in [1]. The methodology is based on multivariate discrete wavelet transform (DWT). Technically, it lies in the computation of the DWT coefficients of each component Y_m of Y, $D_{Y_m}(2^j, k) = \langle 2^{-j/2}\psi_0(2^{-j}t-k)|Y_m(t)\rangle, \forall k \in \mathbb{Z}, \forall j \in \{j_1, \ldots, j_2\}$, with the same mother wavelet ψ_0 [10]. The multivariate DWT of Y results from the concatenation of the DWT coefficients along components, $D_Y(2^j, k) = (D_{Y_1}(2^j, k), \ldots, D_{Y_M}(2^j, k))$. The $M \times M$ covariance matrices of multivariate DWT along scales 2^j, called *wavelet spectrum*, read

$$S(2^j) \triangleq \frac{1}{n_j}\sum_{k=1}^{n_j} D_Y(2^j, k)D_Y(2^j, k)^*, \quad (2)$$

with n_j the number of available wavelet coefficients at scale 2^j. It can be shown that the eigenvalues $\lambda_m(2^j)$, $m = 1, \ldots, M$, of $S(2^j)$ asymptotically follow power laws with exponents $2H_m+1$, naturally leading to estimate H_m by linear regression on the logarithm of $\lambda_m(2^j)$ against scales 2^j, i.e.,

$$\hat{H}_m = \frac{1}{2}\sum_{j=j_1}^{j_2} v_j \log_2 \lambda_m(2^j) - \frac{1}{2}, \ \forall m = 1, \ldots, M, \quad (3)$$

with weights v_j verifying $\sum_j jv_j = 1$ and $\sum_j v_j = 0$.

This procedure is affected by the so-called *repulsion effect* [12], a finite-size scale-dependent bias that emerges when the exponents H_m are close: the estimated eigenvalues $\lambda_m(2^j)$ deviate more than the exact eigenvalues on average at scale 2^j, especially as the number n_j of available wavelet coefficients to compute $S(2^j)$ is small. Since repulsion effects between $\lambda_m(2^j)$ depends on scale 2^j, this results in a bias in the linear regressions to obtain $\hat{H}_m$. A procedure to reduce the impact of this bias in the linear regression, proposed in [9], consists in computing wavelet spectra $S^{(w)}(2^j)$ from several non-overlapping time windows $w = 1, \ldots, 2^{j_2-j}$ of DWT coefficients of same size n_{j_2} along scales 2^j, i.e.,

$$S^{(w)}(2^j) \triangleq \frac{1}{n_{j_2}}\sum_{k=1+(w-1)n_{j_2}}^{wn_{j_2}} D_Y(2^j, k)D_Y(2^j, k)^*, \quad (4)$$

so that the eigenvalues $\lambda_m^{(w)}(2^j)$, $m = 1, \ldots, M$, of $S^{(w)}(2^j)$ are affected by similar repulsion effects across scales 2^j. Finally, the linear regressions are now performed on the logarithms of $\lambda_m^{(w)}(2^j)$ averaged across windows w,

$$\log_2 \bar{\lambda}_m(2^j) \triangleq 2^{j-j_2}\sum_{w=1}^{2^{j_2-j}} \log_2 \lambda_m^{(w)}(2^j), \quad (5)$$

against scales 2^j, as follows:

$$\hat{H}_m^{(\mathrm{bc})} = \frac{1}{2}\sum_{j=j_1}^{j_2} v_j \log_2 \bar{\lambda}_m(2^j) - \frac{1}{2}, \quad \forall m = 1, \ldots, M. \quad (6)$$

See [9] for the empirical assessment of this estimation procedure on synthetic finite-size M-fBm.

3 High-dimensional asymptotics

We consider the following high-dimensional setting. First, the Hurst exponent vector $\underline{H}$ is not deterministic but is a vector of M i.i.d. samples from a discrete distribution $\pi(dH)$ with support $\{H_1, ..., H_L\}$, $0 < H_1 \leq \ldots \leq H_L < 1$, $L \in \mathbb{N}$. Second, we consider that the number of components M is not fixed but goes to infinity as sample size N goes to infinity. In the context of wavelet analysis, it entails to perform the linear regressions to estimate $\underline{H}$ on a range of analysis scales $2^{j_1} \leq 2^j \leq 2^{j_2}$ also going to infinity as $N \to +\infty$ with $j_2 - j_1$ constant. Specifically, the high-dimensional behavior implies a *three-way limit*:

$$\frac{M}{N/2^{j_2}} \to c \in [0, +\infty) \quad \text{as } M, N, j_2 \to +\infty. \quad (7)$$

As a consequence, since the number n_{j_2} of available wavelet coefficients at scale 2^{j_2} approximately behaves as $n_{j_2} \approx N/2^{j_2}$, the ratio M/n_{j_2} is asymptotically constant.

In this setting, it has been shown in [11] that the number of modes in the distribution of the M estimates $\hat{\underline{H}} = (\hat{H}_1, \ldots, \hat{H}_M)$ obtained from Eq. (3) asymptotically tends to the number of modes of $\pi(dH)$ as $N \to +\infty$. We can show that this result still holds for the M estimates $\hat{\underline{H}}^{(\mathrm{bc})} = (\hat{H}_1^{(\mathrm{bc})}, \ldots, \hat{H}_M^{(\mathrm{bc})})$ obtained from Eq. (6). This is assessed numerically in Section 5. In practice, this means that $\pi(dH)$ is unimodal if the distribution of $\hat{\underline{H}}$ is unimodal, and multimodal otherwise.

4 Bootstrap unimodality testing

4.1 Dip test

Since the unimodality of the distribution of $\hat{\underline{H}}^{(\mathrm{bc})} = (\hat{H}_1^{(\mathrm{bc})}, \ldots, \hat{H}_M^{(\mathrm{bc})})$ reproduces the unimodality of the distribution $\pi(dH)$ of the entries of $\underline{H}$, we devise a test procedure based on the distribution of the M estimates $\hat{H}_m^{(\mathrm{bc})}$ to test the null hypothesis

$$\mathcal{H}_0 : \quad \pi(dH) \text{ is unimodal.} \quad (8)$$

There exists several unimodality test procedures. Following [11], we use Hartigan's dip test [6]. Let $\hat{F}$ be the empirical cumulative distribution function of the estimates $\hat{H}_m^{(\mathrm{bc})}$, i.e.,

$$\hat{F}(x) = \frac{1}{M}\sum_{m=1}^{M} \mathbb{1}_{\{\hat{H}_m^{(\mathrm{bc})} < x\}}, \quad \forall x \in \mathbb{R}. \quad (9)$$

Let $\mathcal{U}$ be the class of all unimodal cumulative distribution functions. The dip statistic, defined as

$$\hat{d} = \inf_{G\in\mathcal{U}} \sup_{x\in\mathbb{R}} |\hat{F}(x) - G(x)|, \quad (10)$$

measures the deviation of the empirical cumulative distribution function to a unimodal distribution function. The test procedure reads

$$\text{reject } \mathcal{H}_0 \text{ when } \hat{d} > d_\alpha, \quad (11)$$

with $\alpha \in (0,1)$ the significance level, defined in [8] by $\alpha = \sup_{H_0 \in (0,1)} \mathbb{P}(\hat{d} > d_\alpha : \mathrm{supp}(\pi(dH)) = H_0)$. In [11], d_α is estimated for fixed α using the distribution of $\hat{d}$ under $\mathcal{H}_0$ across Monte Carlo realizations of synthetic ofBm. In the present work, we propose to approximate d_α from a single observation of multivariate data using a time-scale block-bootstrap procedure developed in [9].

4.2 Bootstrap resampling

A multivariate block-bootstrap resampling scheme is performed in the wavelet domain. This approach allows to preserve the time-scale multivariate dependence structure as opposed to a resampling of the components independently [7]. Technically, we perform sampling with replacement of $\lceil n_j / L_B \rceil$ possibly overlapping blocks of multivariate DWT coefficients $(D(2^j, k), \ldots, D(2^j, k + L_B - 1))$, $k = 1, \ldots, n_j$, of size L_B in time, at each scale 2^j to obtain R bootstrap resamples $D_j^{*(r)} = (D^{*(r)}(2^j, 1), \ldots, D^{*(r)}(2^j, n_j))$, $r = 1, \ldots, R$. From each $D_j^{*(r)}$, we successively compute estimates $S^{*(r,w)}(2^j)$, for all windows $w = 1, \ldots, 2^{j_2 - j}$, then $\log_2 \bar{\lambda}_m^{*(r)}(2^j)$ and finally $\hat{H}_m^{*(\mathrm{bc},r)}$ using Eqs. (4-6). For simplification of notations, $\hat{H}_m^{*(\mathrm{bc},r)}$ is replaced by $\hat{H}_m^{*(r)}$. The bootstrap Hurst exponent estimates are then centered to reproduce null hypothesis $\mathcal{H}_0$, $\bar{H}_m^{*(r)} = \hat{H}_m^{*(r)} - \langle \hat{H}_m^* \rangle$, with $\langle \hat{H}_m^* \rangle$ the average of bootstrap estimates over the samples, $\langle \hat{H}_m^* \rangle = 1/R \sum_r \hat{H}_m^{*(r)}$.

For each $r = 1, \ldots, R$, we expect the distribution of the boostrap estimates $\bar{H}_m^{*(r)}$ along components $m = 1, \ldots, M$ to be unimodal. We therefore compute the empirical cumulative distribution function $\hat{F}^{*(r)}$ from estimates $\bar{H}_m^{*(r)}$, $m = 1, \ldots, M$, for each bootstrap resample $r = 1, \ldots, R$ using Eq. (9) and then the resulting dip statistic $\hat{d}^{*(r)}$ as in Eq. (10). Finally, the threshold d_α for the dip test (11) is estimated by $d_\alpha^* = \hat{d}^{*(l)}$ with $l = 1, \ldots, R$ such that

$$\#\{r = 1, \ldots, R : \hat{d}^{*(r)} > \hat{d}^{*(l)}\} = \lfloor (1 - \alpha) R \rfloor, \qquad (12)$$

i.e., $\hat{d}^{*(l)}$ is approximately a α percentage of the bootstrap dip statistics $\hat{d}^{*(r)}$, $r = 1, \ldots, R$. The bootstrap dip test then reads

$$\text{reject } \mathcal{H}_0 \text{ when } \hat{d} > d_\alpha^*. \qquad (13)$$

5 Performance assessment

5.1 Monte Carlo experiment setting

To evaluate the relevance of the bootstrap unimodality test procedure, we conduct numerical experiments on $N_{\mathrm{MC}} = 100$ realizations of synthetic M-fBm. In the context of the three-way limit (7), we consider several numbers of components $M \in \{2^4, 2^5, 2^6\}$ and the linear regressions for the estimation are performed on scales $2^{j_1} \leq 2^j \leq 2^{j_2}$ depending on M for different sample sizes N such that $c \triangleq M\, 2^{j_2}/N$ is fixed. The closer the value of c is to 1, the slower the convergence of the high-dimensional asymptotic behavior of $\underline{\hat{H}}^{(\mathrm{bc})}$ described in Section 3. We therefore consider two values of $c \in \{1/8, 1/4\}$. Moreover, to explore various relations between M, N, j_1 and j_2 for a same c, two ranges of analysis scales are examined at a ratio c fixed: the linear regressions are performed either from $2^{j_1} = M/4$ to $2^{j_2} = M$ or from $2^{j_1} = M/8$ to $2^{j_2} = M/2$. The M entries of the Hurst exponent vector $\underline{H}$ are sampled uniformly from the support $\{H_1, H_2\}$ where $H_1 = 0.6$ and $H_2 = H_1 + \Delta H$. The covariance matrix Σ_X is set to identity matrix and the mixing matrix W is randomly chosen among $M \times M$ orthogonal matrices and kept fixed for each set (M, N). For the estimation procedure, the multivariate DWT is computed using the Daubechies 2 mother wavelet. For the test procedure, $R = 500$ block-bootstrap of size $L_B = 4$ (size of the wavelet support, following [13]) are resampled from the multivariate wavelet coefficients.

5.2 High-dimensional asymptotics for $\underline{\hat{\mathbf{H}}}$

The behavior of the estimates $\hat{H}_m^{(\mathrm{bc})}$, $m = 1, \ldots, M$, of the entries of the Hurst exponent vector $\underline{H}$ is examined in different settings. Fig. 1 reports the histograms of the estimates $\hat{H}_m^{(\mathrm{bc})}$ across entries m and Monte Carlo realizations for three different distribution supports $\{H_1, H_2\}$ of $\underline{H}$ and both values of c with different dimensions M, sample sizes N and analysis octaves (j_1, j_2). These results first show that, when $H_1 \neq H_2$ ($\mathcal{H}_0$ not true), several modes appear in the distribution of $\hat{H}_m^{(\mathrm{bc})}$ as N, M, j_1 and j_2 are jointly increasing at c fixed, corroborating the high-dimensional asymptotic behavior of the M estimates $\hat{H}_m^{(\mathrm{bc})}$ stated in Section 3. Moreover, it can be observed that convergence of the estimated distributions of $\hat{H}_m^{(\mathrm{bc})}$ is faster for $c = 1/8$ compared to $c = 1/4$ for comparable analysis octaves (j_1, j_2) (second and fifth rows) and similar between the two values of c for comparable sample sizes N (third and fifth rows).

5.3 Reproduction of the null hypothesis

To assess the bootstrap procedure, we first study the behavior of the test procedure under the null hypothesis $\mathcal{H}_0$ ($H_1 = H_2$). Fig. 2 reports the decisions $\hat{d} > d_\alpha^*$ of the bootstrap dip test to reject $\mathcal{H}_0$ averaged across Monte Carlo realizations (with 95% confidence interval) as a function of the preset significance level α for different three-way limits c related to different sample sizes N, dimensions M and analysis octaves (j_1, j_2). For both values of c, the targeted significance levels α are well reproduced by the bootstrap procedure.

5.4 Test power

We quantify the power of the test. Fig. 3 relates the empirical power (Monte Carlo average with 95% confidence interval) as a function of the gap $\Delta H = H_2 - H_1$ for a preset significance level $\alpha = 0.05$ and different three-way limits c kept fixed for different sample sizes N, dimensions M and analysis octaves (j_1, j_2). These results first show that, at a fixed ratio c, power is increasing with sample size N. In addition, at a ratio M/N fixed (i.e., comparing left black line with '+' and right blue line with 'Δ', and left blue line with 'Δ' and right red line with 'o'), power is increasing with M, that is for a larger number of samples used to compute the empirical cumulative distribution function (9) involved in computing the dip statistic.

6 Conclusion

The present work designs an original Hurst unimodality test procedure for high-dimensional multivariate data based on Hartigan's dip statistic by combining an eigen-wavelet estimation procedure and a time-scale block-bootstrap resampling

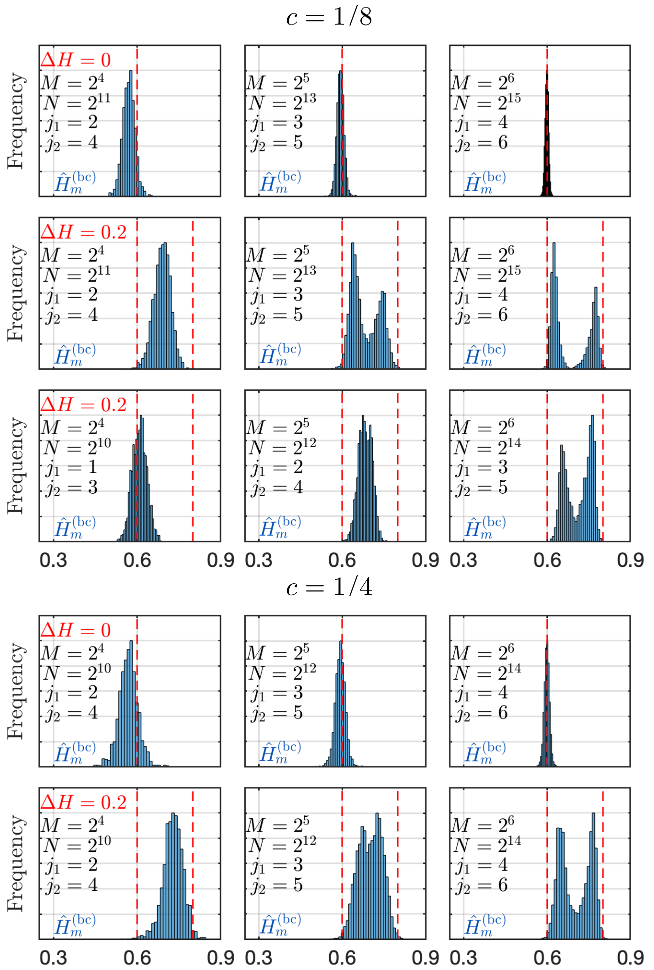

Figure 1 – **Distribution of $\hat{\mathbf{H}}_{\mathbf{m}}^{(\mathrm{bc})}$ in the high-dimensional three-way limit.** Histograms of $\hat{H}_m^{(\mathrm{bc})}$ over Monte Carlo realizations and components $m = 1, \dots, M$ for different dimensions M, sample sizes N and analysis octaves (j_1, j_2) such that (first three lines) $c = 1/8$ and (last two lines) $c = 1/4$, and for (dashed red lines) different values H_1 and H_2 of the distribution support $\{H_1, H_2\}$ of $\underline{H}$.

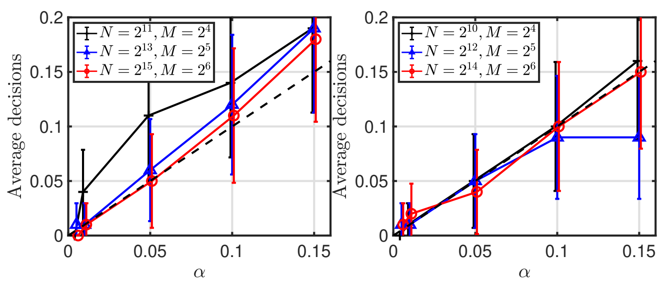

Figure 2 – **Significance levels.** Rejection decisions of the test averaged over Monte Carlo realizations (with 95% confidence interval) against the preset significance level α under $\mathcal{H}_0$ for different dimensions M, sample sizes N and analysis octaves (j_1, j_2) such that (left) $c = 1/8$ and (right) $c = 1/4$.

scheme. The relevance and performance of the test procedure are assessed by Monte Carlo simulations of synthetic finite-size M-fBm. Results first show that the distribution of the Hurst exponent vector estimator is unimodal or multimodal depending on the presence of a single or multiple Hurst exponents driving the M-fBm asymptotically, in the context of a high-dimensional three-way limit as both dimension and sample size increase. Results also confirm that the bootstrap

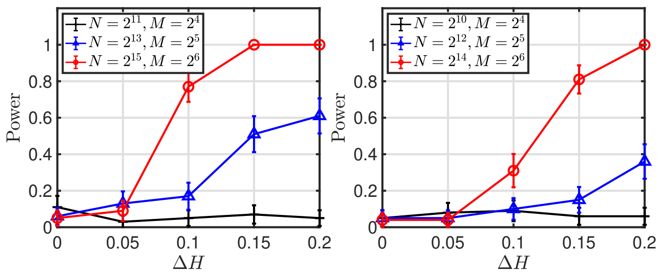

Figure 3 – **Test power.** Proportions of rejections (Monte Carlo average with 95% confidence interval) against the gap $\Delta H = H_2 - H_1$ between Hurst exponents of the distribution support $\{H_1, H_2\}$ of $\underline{H}$ for a preset significance level $\alpha = 0.05$ and different dimensions M, sample sizes N and analysis octaves (j_1, j_2) such that (left) $c = 1/8$ and (right) $c = 1/4$.

procedure is powerful asymptotically, that is for realistic sample sizes N and large dimensions M. Future work could include the identification of the number of modes in the Hurst exponent vector when unimodality has been rejected. Matlab routines will be made available at the time of publication.

References

[1] P. Abry and G. Didier. Wavelet eigenvalue regression for n-variate operator fractional Brownian motion. *J. Multivar. Anal.*, 168:75–104, November 2018.

[2] P.-O. Amblard and J.-F. Coeurjolly. Identification of the multivariate fractional Brownian motion. *IEEE Trans. Signal Process.*, 59(11):5152–5168, 2011.

[3] Z. Bai and J. W. Silverstein. *Spectral Analysis of Large Dimensional Random Matrices*, volume 20. Springer, New York, 2nd edition, 2010.

[4] J.-F. Coeurjolly, P.-O. Amblard, and S. Achard. Wavelet analysis of the multivariate fractional Brownian motion. *ESAIM Prob. Stat.*, 17:592–604, 2013.

[5] G. Didier and V. Pipiras. Integral representations and properties of operator fractional Brownian motions. *Bernoulli*, 17(1):1–33, 2011.

[6] J. A. Hartigan and P. M. Hartigan. The dip test of unimodality. *Ann. Stat.*, pages 70–84, 1985.

[7] S. N. Lahiri. *Resampling Methods for Dependent Data*. Springer, New York, 2003.

[8] E. L. Lehmann, J. P. Romano, and G. Casella. *Testing statistical hypotheses*, volume 3. Springer, 2005.

[9] C.-G. Lucas, P. Abry, H. Wendt, and G. Didier. Bootstrap for testing the equality of selfsimilarity exponents across multivariate time series. In *Proc. IEEE Eur. Signal Process. Conf. (EUSIPCO)*, pages 1960–1964, 2021.

[10] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, CA, 1998.

[11] O. Orejola, G. Didier, P. Abry, and H. Wendt. Hurst multimodality detection based on large wavelet random matrices. In *Proc. IEEE Eur. Signal Process. Conf. (EUSIPCO)*, pages 2131–2135, 2022.

[12] T. Tao. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.

[13] H. Wendt, P. Abry, and S Jaffard. Bootstrap for empirical multifractal analysis. *IEEE Signal Proces. Mag.*, 24(4):38–48, 2007.