Decomposing Spectral and Phasic Differences in Nonlinear Features between Datasets

Pedro A. M. Mediano^{1,*}, Fernando E. Rosas^{2,3,4,†}, Adam B. Barrett^{5,‡} and Daniel Bor^{1,§}¹Department of Psychology, University of Cambridge, Cambridge CB2 3EB, United Kingdom²Centre for Psychedelic Research, Department of Medicine, Imperial College London, London SW7 2DD, United Kingdom³Data Science Institute, Imperial College London, London SW7 2AZ, United Kingdom⁴Centre for Complexity Science, Imperial College London, London SW7 2AZ, United Kingdom⁵Sackler Centre for Consciousness Science, Department of Informatics, University of Sussex, Brighton BN1 9RH, United Kingdom

(Received 13 February 2020; revised 17 May 2021; accepted 15 July 2021; published 17 September 2021)

When employing nonlinear methods to characterize complex systems, it is important to determine to what extent they are capturing genuine nonlinear phenomena that could not be assessed by simpler spectral methods. Specifically, we are concerned with the problem of quantifying spectral and phasic effects on an observed difference in a nonlinear feature between two systems (or two states of the same system). Here we derive, from a sequence of null models, a decomposition of the difference in an observable into spectral, phasic, and spectrum-phase interaction components. Our approach makes no assumptions about the structure of the data and adds nuance to a wide range of time series analyses.

DOI: [10.1103/PhysRevLett.127.124101](https://doi.org/10.1103/PhysRevLett.127.124101)

Nonlinear methods are useful for characterizing differences between various states of a complex system, and have found applications in a wide range of scientific domains. For example, Lempel-Ziv (LZ) complexity [1] and multiscale entropy [2] have been successful in discriminating between conscious and unconscious brain activity [3], and have yielded insights into physiological pathologies [4] and price dynamics [5]. However, more refined conclusions could be obtained if there were a principled way to assess how much of the differences in such measures are due to genuine nonlinear effects, and how much is explainable by changes in the power spectrum.

A popular approach to study the effect of spectral and phasic contributions on an observable is via surrogate data methods [6], which examine whether its value is representative of a null distribution obtained from surrogate data. Such surrogate methods are regarded as a basic constituent of the data analyst's toolkit [7], and have been extended to a range of scenarios including multivariate time series [8], nonstationary data [9], and many others. However, surrogate methods are typically designed to be applied on a *single dataset*, and it is not straightforward to use them to disentangle spectral and phasic contributions on differences in an observable *between two datasets*—e.g., how much of

the difference in LZ complexity between two neurological conditions simply reflects the known spectral changes between them [10]. The crux of why this is challenging, and why naive applications of typical surrogate methods fail, is that the difference between two null models is not necessarily a good null model of the difference (see Supplemental Material [11] for a detailed example).

To deal with this issue, here we present a novel decomposition of the difference in an observable between two time series datasets into spectral, phasic, and spectrum-phase interaction components. The decomposition makes no assumptions about the structure of the data, and is widely applicable to a broad range of scenarios of interest. We illustrate our method by analyzing LZ complexity on neuroimaging data, where our decomposition identifies phasic and spectrum-phase interaction components that take the opposite sign to the predominantly spectral overall effect, and which would not have been detectable by previously existing methods.

The decomposition.—Let us consider a scientist who is interested in an observed difference in some quantity f between data recorded in two different conditions, denoted by $\mathcal{X}$ and $\mathcal{Y}$. The data consist of time series recordings, and a set of time series segments are obtained from each condition. Each segment could correspond to data recorded from, e.g., different participants in an experiment, or different time periods from the same participant. The whole dataset from the first condition is denoted as $\mathbf{x}^N$, where N is the population size of these data, and the N time series segments within $\mathbf{x}^N$ as $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$. Similarly, for the second condition one has $\mathbf{y}^M = \{\mathbf{y}_1, \dots, \mathbf{y}_M\}$. Our goal is to decompose the difference in f between $\mathcal{X}$ and $\mathcal{Y}$

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

into spectral, phasic, and spectrum-phase interaction components—i.e., to decompose

$$\Delta(\mathbf{x}^N, \mathbf{y}^M) := \bar{f}(\mathbf{x}^N) - \bar{f}(\mathbf{y}^M), \quad (1)$$

where $\bar{f}(\mathbf{x}^N) = (1/N) \sum_{j=1}^N f(\mathbf{x}_j)$ and $\bar{f}(\mathbf{y}^M) = (1/M) \sum_{k=1}^M f(\mathbf{y}_k)$ are the empirical ensemble averages of the function in question, f . This is achieved by a series of comparisons between expected f values on the data and those on a set of progressively more constrained null models for the stochastic processes underlying the data.

Formally, we consider $\mathbf{x}_1, \dots, \mathbf{x}_N$ to be independent and identically distributed (i.i.d.) realizations of a stochastic process sampled under condition $\mathcal{X}$, and $\mathbf{y}_1, \dots, \mathbf{y}_M$ to be i.i.d. realizations of another stochastic process sampled under condition $\mathcal{Y}$, and $\mathbf{x}_j, \mathbf{y}_k \in \mathbb{R}^T$, where T is the length of each time series. The decomposition utilizes the discrete Fourier transform, which is denoted by $\hat{\mathbf{x}} = \mathcal{F}\{\mathbf{x}\} \in \mathbb{C}^T$, given a time series $\mathbf{x}$. The amplitudes of the Fourier components are denoted by $\mathbf{A}(\hat{\mathbf{x}}) = \{A_1(\hat{\mathbf{x}}), \dots, A_T(\hat{\mathbf{x}})\} \in \mathbb{R}^T$, and their phases by $\boldsymbol{\phi}(\hat{\mathbf{x}}) = \{\phi_1(\hat{\mathbf{x}}), \dots, \phi_T(\hat{\mathbf{x}})\} \in [0, 2\pi]^T$. Thus, the data for $\mathcal{X}$ can be represented in the frequency domain as i.i.d. phase-amplitude tuples $(\mathbf{A}(\hat{\mathbf{x}}_j), \boldsymbol{\phi}(\hat{\mathbf{x}}_j))$, following a distribution $p_{\mathcal{X}}(\mathbf{A}, \boldsymbol{\phi})$ induced by $\mathcal{X}$ —and similarly for the $\mathbf{y}_k$.

We begin by considering a null model $\mathcal{M}^i$ on which amplitudes and phases have no interaction—i.e., are statistically independent. Accordingly, we construct new time series $\mathbf{x}_j^{(w)}$ that satisfy this null model by combining the spectrum of each $\mathbf{x}_j$ with the phases from some other randomly chosen time series from *within* condition $\mathcal{X}$ (and similarly for the $\mathbf{y}_k$). That is, we construct $\mathbf{x}_j^{(w)} = \mathcal{F}^{-1}\{\mathbf{A}(\hat{\mathbf{x}}_j)e^{i\phi(\hat{\mathbf{x}}_{\alpha_j})}\}$ and $\mathbf{y}_k^{(w)} = \mathcal{F}^{-1}\{\mathbf{A}(\hat{\mathbf{y}}_k)e^{i\phi(\hat{\mathbf{y}}_{\beta_k})}\}$, where α_j and β_k are distributed uniformly over $\{1, \dots, N\}$ and $\{1, \dots, M\}$, respectively. We then consider the mean value of f on these phase-shuffled data, given by $\nu^i(\mathbf{x}^N) := (1/N) \sum_{j=1}^N f(\mathbf{x}_j^{(w)})$. The spectrum-phase interaction contribution to the value of f in condition $\mathcal{X}$ is then calculated as

$$\Delta^i(\mathbf{x}^N) := \bar{f}(\mathbf{x}^N) - \mathbb{E}\{\nu^i(\mathbf{x}^N)|\mathbf{x}^N\}, \quad (2)$$

where the conditional expectation averages the effect of the random integers α_j on ν^i . Similarly, $\Delta^i(\mathbf{y}^M)$ can be calculated for $\mathcal{Y}$. When estimating $\Delta^i(\mathbf{x}^N)$ and $\Delta^i(\mathbf{y}^M)$ in practice, one will approximate the distribution of ν^i by averaging multiple realizations of it.

The quantity $\Delta^i(\mathbf{x}^N)$ measures the extent to which the expected value of f would be affected if one were to break any dependence that exists between the amplitudes and phases of the $\hat{\mathbf{x}}_j$. Equivalently, $\Delta^i(\mathbf{x}^N)$ accounts for the deviation in the mean value of f in condition $\mathcal{X}$ from that which would be expected if the null model $\mathcal{M}^i$ holds. For large N , the law of large numbers guarantees that

$$\Delta^i(\mathbf{x}^N) \rightarrow \mathbb{E}_{p_{\mathcal{X}}(\mathbf{A}, \boldsymbol{\phi})}\{f(\mathbf{A}, \boldsymbol{\phi})\} - \mathbb{E}_{p_{\mathcal{X}}(\mathbf{A})p_{\mathcal{X}}(\boldsymbol{\phi})}\{f(\mathbf{A}, \boldsymbol{\phi})\},$$

and hence that in the absence of any dependency between the phases and spectra, i.e., when $p_{\mathcal{X}}(\mathbf{A}, \boldsymbol{\phi}) = p_{\mathcal{X}}(\mathbf{A})p_{\mathcal{X}}(\boldsymbol{\phi})$, $\lim_{N \rightarrow \infty} \Delta^i(\mathbf{x}^N) = 0$.

Next, we focus on the *phasic effect* on f , i.e., the effect of differences between the phase distributions of $\mathcal{X}$ and $\mathcal{Y}$. For this, we consider a second null model $\mathcal{M}^\phi$ under which phases are not only independent from amplitude but also follow the same distribution in each of the conditions $\mathcal{X}$ and $\mathcal{Y}$. We construct phase-shuffled time series $\mathbf{x}_j^{(a)}, \mathbf{y}_k^{(a)}$ that satisfy this null model by replacing the phases of each time series with those from another randomly chosen time series from the whole set of data $\{\mathbf{x}^N, \mathbf{y}^M\}$. That is, we construct $\mathbf{x}_j^{(a)} = \mathcal{F}^{-1}\{\mathbf{A}(\hat{\mathbf{x}}_j)e^{i\phi(\hat{\mathbf{w}}_j)}\}$ and $\mathbf{y}_k^{(a)} = \mathcal{F}^{-1}\{\mathbf{A}(\hat{\mathbf{y}}_k)e^{i\phi(\hat{\mathbf{z}}_k)}\}$, where $\hat{\mathbf{w}}_j, \hat{\mathbf{z}}_k$ are the discrete Fourier transforms of independently randomly chosen time series that are each drawn from $\mathcal{X}$ with probability 1/2, and from $\mathcal{Y}$ with probability 1/2. Then, we consider the mean value of f on these phase-shuffled data: $\nu^\phi(\mathbf{x}^N|\mathbf{y}^M) := (1/N) \sum_{j=1}^N f(\mathbf{x}_j^{(a)})$, and introduce

$$\Delta^\phi(\mathbf{x}^N) := \mathbb{E}\{\nu^i(\mathbf{x}^N) - \nu^\phi(\mathbf{x}^N|\mathbf{y}^M)|\mathbf{x}^N, \mathbf{y}^M\}.$$

We define $\Delta^\phi(\mathbf{y}^M)$ analogously. Again, when estimating these quantities in practice, one can approximate the distributions of ν^i and ν^ϕ , for each condition, by averaging multiple realizations of them.

The quantity $\Delta^\phi(\mathbf{x}^N)$ measures the expected effect on the mean value of f in condition $\mathcal{X}$ if $\mathcal{M}^\phi$ holds—i.e., the effect of changing the probability distribution of the phases from $p_{\mathcal{X}}(\boldsymbol{\phi})$ to the mixture $[p_{\mathcal{X}}(\boldsymbol{\phi}) + p_{\mathcal{Y}}(\boldsymbol{\phi})]/2$. Note that if the distribution of phases is the same for both conditions, so that $p_{\mathcal{X}}(\boldsymbol{\phi}) = p_{\mathcal{Y}}(\boldsymbol{\phi})$, then the law of large numbers guarantees that

$$\lim_{N \rightarrow \infty} \Delta^\phi(\mathbf{x}^N) = \lim_{M \rightarrow \infty} \Delta^\phi(\mathbf{y}^M) = 0. \quad (3)$$

Finally, we consider the effect of spectral differences between the conditions on the difference in f . For this, we consider the deviation of the phase-shuffled data above from a further constrained null model $\mathcal{M}^A$, in which both amplitudes and phases are statistically independent and distributed identically in $\mathcal{X}$ and $\mathcal{Y}$. Specifically, we consider $\nu^A(\mathbf{x}^N, \mathbf{y}^M) := \nu^\phi(\mathbf{x}^N|\mathbf{y}^M) - \nu^\phi(\mathbf{y}^M|\mathbf{x}^N)$. Since $\mathbf{x}_i^{(a)}$ and $\mathbf{y}_i^{(a)}$ have, by definition, the same phase statistics, $\nu^\phi(\mathbf{x}^N|\mathbf{y}^M)$ and $\nu^\phi(\mathbf{y}^M|\mathbf{x}^N)$ will, on average, differ only because of differences between the distribution of the spectrum of $\mathcal{X}$ and $\mathcal{Y}$. Therefore, we introduce

$$\Delta^A(\mathbf{x}^N, \mathbf{y}^M) := \mathbb{E}\{\nu^A(\mathbf{x}^N, \mathbf{y}^M)|\mathbf{x}^N, \mathbf{y}^M\}$$

as a metric of the spectral effect. If the distribution of the spectrum is the same for both conditions, then

$\lim_{N,M \rightarrow \infty} \Delta^A(\mathbf{x}^N, \mathbf{y}^M) = 0$ by the law of large numbers. Again, when estimating this quantity in practice, one will approximate the distribution of $\mathcal{L}^A(\mathbf{x}^N, \mathbf{y}^M)$ by obtaining multiple realizations.

With these quantities at hand, via a telescopic sum we can obtain a decomposition of the total difference in f between the two conditions into spectral, phasic, and spectrum-phase interaction terms. We have that the difference in mean f values between the conditions decomposes into

$$\begin{aligned} \Delta(\mathbf{x}^N, \mathbf{y}^M) = & \Delta^A(\mathbf{x}^N, \mathbf{y}^M) \quad \text{Spectrum} \\ & + \Delta^\phi(\mathbf{x}^N) - \Delta^\phi(\mathbf{y}^M) \quad \text{Phase} \\ & + \Delta^i(\mathbf{x}^N) - \Delta^i(\mathbf{y}^M) \quad \text{Interaction.} \end{aligned} \quad (4)$$

Of these, the first term is the difference in f that persists on data modified so the phases have the same distribution across conditions, and so corresponds to the difference attributable to spectral changes only. Similarly, by comparing the data with the observed phase distributions against data with identically distributed phases, the second term measures the difference in f attributable to phase changes. Finally, the third term compares the observed data with phase-shuffled time series to account for changes due to the phase-spectrum interaction in both conditions.

Accordingly, each of the Δ 's can be considered to be comparing expected f values on the data against f

values on a set of increasingly restrictive null models, $\mathcal{M}^i \rightarrow \mathcal{M}^\phi \rightarrow \mathcal{M}^A$. We note that this decomposition is invariant to the order in which the decomposition is constructed, i.e., it doesn't make a difference if phasic effects are considered before spectral contributions (as described here), or vice versa (proof in Supplemental Material).

Example.—As an illustration, we present an analysis of the entropy rate of binarized magnetoencephalographic (MEG) signals, as measured with LZ complexity. We use the Cambridge Centre for Ageing and Neuroscience (CAMCAN) dataset [12], which includes a large-scale MEG dataset of participants undergoing several cognitive tasks, and study the differences in Lempel-Ziv complexity [1] between participants in wakeful rest, and participants performing a simple cognitive stop-signal, go/no-go task [12]. This measure (or minor variations of it) has been widely used in the neuroscience literature [13–16], showing a remarkable performance in discriminating between different states of consciousness, for instance normal wakefulness versus sleep [3].

In this application, we consider data for 131 participants in both “task” and “rest” conditions. The data from each participant were divided up into 100 nonoverlapping windows of length $T = 1024$ (which corresponds to approximately 4 s given the sampling rate of 250 Hz). To compute the LZ complexity, time series were binarized, and then the original (1976) version of the LZ complexity

FIG. 1. Decomposition of the difference in Lempel-Ziv (LZ) complexity between task and rest conditions in the CAMCAN MEG dataset. The LZ complexity computed on sensor-level data during task minus that during rest indicates a pronounced reduction in complexity in frontal areas during task (a). This effect is mostly driven by spectral changes, $\Delta^A(\mathbf{x}^N, \mathbf{y}^M)$ (b). Nevertheless, the decomposition [Eq. (4)] also reveals substantial differences in phase and phase-amplitude interaction contributions between conditions (c). These have a different spatial profile and show the opposite trend to the spectral component [top row shows $\Delta^\phi(\mathbf{x}^N)$ and $\Delta^i(\mathbf{x}^N)$ and bottom row shows $\Delta^\phi(\mathbf{y}^M)$ and $\Delta^i(\mathbf{y}^M)$]. Each quantity is significantly nonzero for most channels (i.e., t-test across participants gives $p \ll 0.05$).

described in Ref. [1] was computed. Binarization was carried out based on the mean value of the time series in question, so the binarized time series contained ones where the raw value was greater than the mean, and zeros where the raw value was less than the mean.

For each of the 204 MEG channels of each participant, the decomposition in Eq. (4) was applied considering $\mathbf{x}^N$ to be the windowed data during task and $\mathbf{y}^M$ to be the windowed data during rest, and using 500 realizations of the random variables involved (i.e., 500 random phase shufflings). Thus, a set of Δ 's was obtained for each channel, for each participant. Then, to assess whether differences were significant at the group level, one-sample t-tests were carried out across participants—for each of the Δ 's, for each channel. The mean value of each of the Δ 's at each MEG channel is shown in Fig. 1.

Our decomposition reveals information about the relation between task and rest that is not captured by other statistical tools. First, by studying the direct difference between LZ complexity in task versus rest, our results show a reduction of complexity in frontal regions, and an increase in the rest of the brain during the task [Fig. 1(a)]. Our decomposition shows that the vast majority of this difference (approximately 7.5 out of 8 units) can be explained by spectral effects [Fig. 1(b)]. Interestingly, contrasting effects are found in the phase and interaction components. In particular, during task there is a strong and heavily localized phase-amplitude interaction component, which becomes much weaker and spatially homogeneous during rest [Fig. 1(c)]. Interestingly, both of these show the *opposite* trend from the direct difference, with an increase in frontal regions and reductions elsewhere during task. The neurobiological implications of these findings will be developed in a separate publication.

Conclusion.—In this Letter we have tackled the problem of determining to what extent a measured difference in some quantity between two time series datasets can be attributed to differences between their power spectra. For this, we introduced a decomposition that uses a sequence of null models to disentangle the effect of spectral, phasic, and phase-amplitude interaction effects. Our decomposition requires no assumptions on the data (beyond that distinct samples within the data are independent), and is easy to compute. As a proof of concept, we provided an example of the decomposition yielding novel results on some neuroimaging data, more nuanced than what was previously possible with a standard analysis of LZ complexity.

Since this decomposition can be applied to any observed difference between two datasets, it promises to be a valuable tool for practitioners in multiple scientific disciplines. Moreover, it will help to deepen our understanding of the behavior of nonlinear properties on datasets describing complex systems.

The authors thank Lionel Barnett and Anil Seth for valuable discussions, and two anonymous referees for comments on earlier versions of this manuscript. We also thank Aleksi Ikkala and Darren Price for vital background work, and Yike Guo for supporting this research. This research was funded in whole, or in part, by the Wellcome Trust (Grant No. 210920/Z/18/Z). F. R. is supported by the Ad Astra Chandaria foundation.

P. M. and F. R. contributed equally to this work. D. B. conceptualized the work. A. B. B. guided the writing of the Letter.

*pam83@cam.ac.uk

†f.rosas@imperial.ac.uk

‡A. B. B. and D. B. are joint senior authors.

- [1] A. Lempel and J. Ziv, *IEEE Trans. Inf. Theory* **22**, 75 (1976).
- [2] M. Costa, A. L. Goldberger, and C.-K. Peng, *Phys. Rev. Lett.* **89**, 068102 (2002).
- [3] A. G. Casali, O. Gosseries, M. Rosanova, M. Boly, S. Sarasso, K. R. Casali, S. Casarotto, M.-A. Bruno, S. Laureys, G. Tononi, and M. Massimini, *Sci. Transl. Med.* **5**, 198ra105 (2013).
- [4] M. Costa and J. Healey, in *Computers in Cardiology 2003* (IEEE, New York, 2003), pp. 705–708.
- [5] E. Martina, E. Rodriguez, R. Escarela-Perez, and J. Alvarez-Ramirez, *Energy Econ.* **33**, 936 (2011).
- [6] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer, *Physica (Amsterdam)* **58D**, 77 (1992).
- [7] G. Lancaster, D. Iatsenko, A. Pidde, V. Ticcinelli, and A. Stefanovska, *Phys. Rep.* **748**, 1 (2018).
- [8] D. Prichard and J. Theiler, *Phys. Rev. Lett.* **73**, 951 (1994).
- [9] J. H. Lucio, R. Valdés, and L. R. Rodríguez, *Phys. Rev. E* **85**, 056202 (2012).
- [10] N. D. Schiff, T. Nauevel, and J. Victor, *Curr. Opin. Neurobiol.* **25**, 7 (2014).
- [11] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.127.124101> for mathematical proofs and additional examples.
- [12] M. A. Shafto, L. K. Tyler, M. Dixon, J. R. Taylor, J. B. Rowe, R. Cusack, A. J. Calder, W. D. Marslen-Wilson, J. Duncan *et al.*, *BMC Neurol.* **14**, 204 (2014).
- [13] X.-S. Zhang, R. J. Roy, and E. W. Jensen, *IEEE Trans. Biomed. Eng.* **48**, 1424 (2001).
- [14] M. M. Schartner, A. Pigorini, S. A. Gibbs, G. Arnulfo, S. Sarasso, L. Barnett, L. Nobili, M. Massimini *et al.*, *Neurosci. Conscious.* **2017**, niw022 (2017).
- [15] D. Dolan, H. J. Jensen, P. Mediano, M. Molina-Solana, H. Rajpal, F. Rosas, and J. A. Sloboda, *Front. Psychol.* **9**, 1341 (2018).
- [16] In these references, all channels are concatenated before the computation. The version used here is what is commonly referred to as “LZs” in the neuroscience literature.