




Detecting frequency modulation in stochastic time-series data

Adrian L. Hauber ^{1,2,*}, Christian Sigloch ^{3,4} and Jens Timmer ^{1,2,5}

¹*Institute of Physics, University of Freiburg, 79104 Freiburg, Germany*

²*Freiburg Center for Data Analysis and Modeling (FDM), University of Freiburg, 79104 Freiburg, Germany*

³*Developmental Biology, Institute Biology I, Faculty of Biology, University of Freiburg, 79104 Freiburg, Germany*

⁴*Spemann Graduate School of Biology and Medicine (SGBM), University of Freiburg, 79104 Freiburg, Germany*

⁵*Centre for Integrative Biological Signalling Studies (CIBSS), University of Freiburg, 79104 Freiburg, Germany*



(Received 1 December 2021; accepted 22 July 2022; published 4 August 2022)

We propose a statistical test to identify nonstationary frequency-modulated stochastic processes from time-series data. Our method uses the instantaneous phase as a discriminatory statistics with reliable critical values derived from surrogate data. We simulated an oscillatory second-order autoregressive process to evaluate the size and power of the test. We found that the test we propose is able to correctly identify more than 99% of nonstationary data when the frequency of the simulated data is doubled after the first half of the time series. Our method is easily interpretable, computationally cheap, and does not require choosing hyperparameters that are dependent on the data.

DOI: [10.1103/PhysRevE.106.024204](https://doi.org/10.1103/PhysRevE.106.024204)

I. INTRODUCTION

Methods to acquire biology data on the cellular level are becoming increasingly relevant [1–4]. Their interpretation, however, may be difficult because data captured at the level of a single cell frequently are of a stochastic nature. This is especially important for time-resolved data where repeating the measurement is often not feasible either from a technical or biological point of view.

Oscillations in the concentration of cellular components such as proteins have been shown to be of great importance, e.g., in the development of organs such as the brain [5]. Identification and analysis of such oscillations is therefore a crucial part of understanding regulatory biological systems that govern development.

Especially the frequency of oscillations is an important regulator for many processes, e.g., the development of vertebrae that is tuned via delta-notch intercellular signaling [6,7]. Given oscillatory data, the question arises whether oscillations have a constant or modulated frequency, i.e., whether the underlying stochastic process is stationary or nonstationary. Here, we propose a method that is specifically tailored to detect frequency modulations in short time series in contrast to existing approaches for identifying nonstationarity [8–14].

We employ the Hilbert transform to obtain instantaneous phase information from a time series to track the speed with which oscillations proceed. We compare this to the instantaneous phase of surrogate time series with constant frequency to derive confidence bands and to identify frequency modulations.

To investigate the size and power of this test, we simulated an oscillating second-order autoregressive process with

different degrees of frequency modulation. We applied our proposed method and counted how often the test identifies a time series as nonstationary depending on the actual frequency modulation.

We found that the test we propose is able to correctly identify more than 99% (97%) of the frequency-modulated time series when the frequency is increased by 100% (80%) with respect to the base frequency.

We also applied our method to a time series from a biological experiment in order to identify the frequency-modulated time series.

II. THE TEST

A. Instantaneous phase as a measure for stationarity

Nonstationarity can for example be exhibited by a time-dependent frequency of oscillation. The instantaneous phase allows us to trace oscillations and thereby detect changes in the frequency, i.e., the speed with which oscillations proceed.

A time series can be thought of as a realization of a stochastic process. Stationary stochastic processes have distribution functions that are shift-invariant in time. The second-order autoregressive [AR(2)] process with variance σ^2 ,

$$x(t) = a_1x(t-1) + a_2x(t-2) + \epsilon(t), \quad \epsilon \sim \mathcal{N}(0, \sigma^2), \quad (1)$$

is an example of a stationary linear stochastic process that can oscillate. The parameters a_1, a_2 can be calculated from the period $T = 1/\nu$ and the mixing time τ [15], which is the decay time of the autocorrelation function, by

$$a_1 = 2 \cos(2\pi\nu)e^{-1/\tau}, \quad (2)$$

$$a_2 = -e^{-2/\tau}. \quad (3)$$

*Corresponding author: adrian.hauber@fdm.uni-freiburg.de

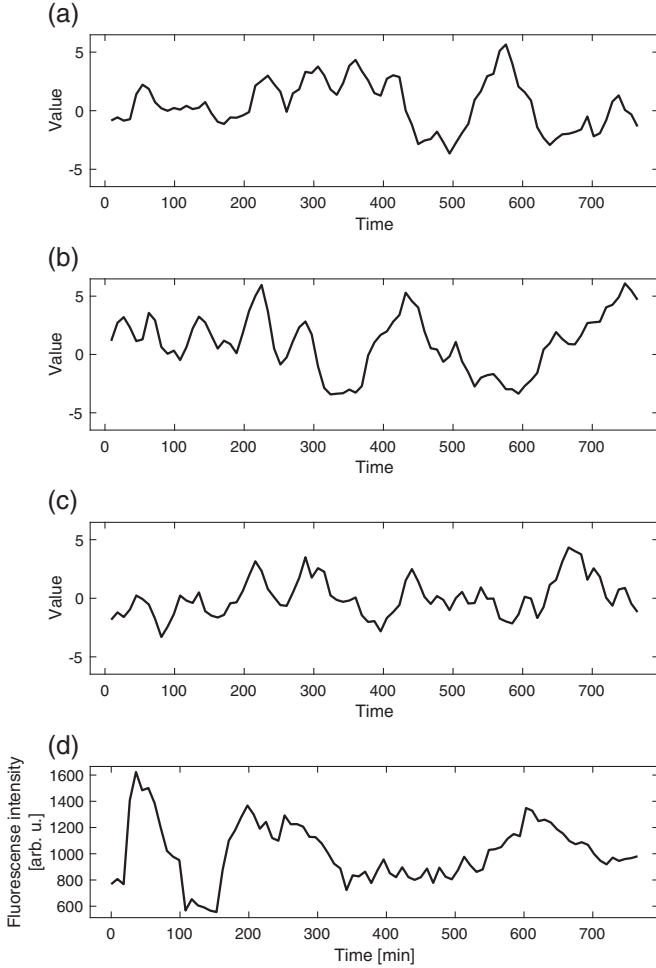


FIG. 1. Three realizations of an oscillating second-order autoregressive process [Eq. (1)] with $\sigma^2 = 1$ and (a) $T = 270$, (b) $T = 180$, (c) $T = 90$, as well as $\tau = 22.5$ and a sampling interval of 9 corresponding to the experimental data in (d). (d) Experimental time series of HES/Her family protein expression in neural stem cells of zebrafish larvae, captured using light sheet microscopy.

Given simulations of this stationary stochastic process [Figs. 1(a)–1(c)], it is hard to decide by visual inspection whether an experimental time series [Fig. 1(d)] is a realization of a stationary process.

Nonstationarity can appear in different manifestations, e.g., a time-dependent frequency which we focus on here. We observed this property in the time series of fluorescence intensities of HES/Her family protein expression in neural stem cells located in the thalamic proliferation zone of zebrafish larvae. It is important to note however that there are many more properties that can change over time due to the underlying stochastic process being nonstationary. For the sake of simplicity, we will skip the reference to the underlying stochastic process and denote the time series itself as “stationary” or “nonstationary.”

Frequency ν is usually defined as the number of oscillations per unit time. Van der Pol however defined frequency as the derivative of a quantity called the instantaneous phase




FIG. 2. (a) Example nonstationary time series with a constant high frequency for $t < 300$ and a constant low frequency for $t \geq 300$. (b) Unwrapped instantaneous phase of the time series in (a). The kink in the instantaneous phase is a consequence of the nonstationarity.

$\Phi(t)$ [16],

$$\nu = \frac{1}{2\pi} \frac{d\Phi(t)}{dt}, \quad (4)$$

which allows us to extend the concept of frequency to time series that have a time-dependent frequency or are “frequency modulated.” For example, a harmonic oscillator $x(t) = A \cos(2\pi\nu t + \psi)$ with $A, \nu, \psi = \text{const}$ has an instantaneous phase of $\Phi(t) = 2\pi\nu t + \psi$, which is a linear function of time.

For a stochastic process that depicts stationary oscillations, the instantaneous phase is also approximately linear in time [Fig. 2(b) for $0 < t < 300$]. When frequency is modulated e.g., with a step function, the instantaneous phase will deviate from the linear appearance and have a kink at the point in time where the adjustment takes place [Fig. 2(b) at $t = 300$].

For a given time series $x(t)$, the instantaneous phase can be calculated as the complex argument of the analytic signal. The analytic signal $x^{(a)}(t) = x(t) + i\tilde{x}(t)$ provides a complex representation of $x(t)$ [17]. Its imaginary part $\tilde{x}(t)$ is calculated by imposing a phase shift of $\pi/2$ to the real-valued time series, which is achieved by applying a discrete Hilbert transform. A computationally cheap way to calculate this transform is to apply a Fourier transform, multiply the positive frequency components by two, neglect the negative frequency components, and to transform back into the time domain [18].

B. Stationary surrogate data serve as control

We use simulated surrogate time series as a reference for interpreting the linearity of the instantaneous phase. These surrogate time series are constituted by random variables that share certain characteristics with the test time series and are stationary by construction. Therefore, a comparison of the surrogate and test time series can reveal nonstationarity of the latter.

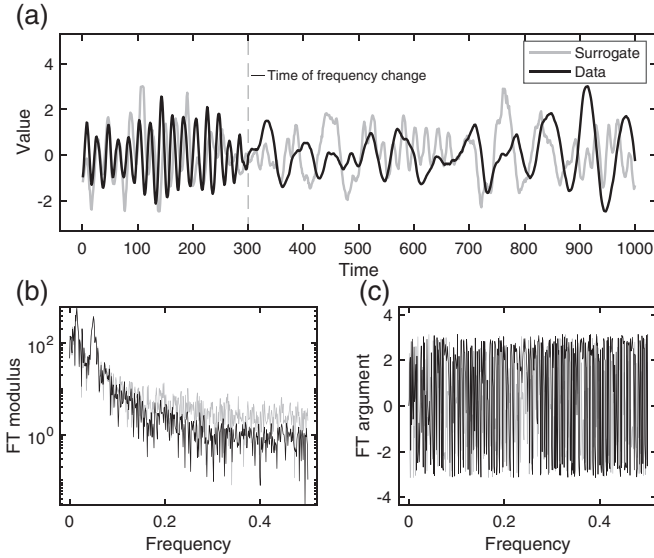


FIG. 3. (a) Example time series (black) from Fig. 2(a) and a surrogate time series (gray). (b) Modulus of the Fourier transform of the two time series in (a). (c) Argument of the Fourier transform of the two time series in (a).

The method of surrogate data [19] is a procedure to test data for a feature by comparing it to simulated data that share most characteristics with the test data but do not include the feature tested for. In our case, we simulate surrogate time series [Fig. 3(a)] from which we construct the time-dependent distribution of instantaneous phases of a stationary stochastic process that approximately share the Fourier transform modulus [Fig. 3(b)] and exactly share the amplitude distribution with the test time series but have randomized Fourier transform arguments [Fig. 3(c)].

To that end, we use the amplitude adjustment Fourier transformation (AAFT) algorithm [19] to generate surrogate time series from a test time series. This algorithm works by reordering random numbers drawn from a normal distribution in such a way that their ranks match those of the test time series. Next, the arguments of this new time series' Fourier transform are randomized. We then back-transform the new time series into the time domain and use it as a reference for reordering the test time series to obtain the surrogate time series. Thereby, the null hypothesis is that of a Gaussian stationary linear process observed via an instantaneous measurement function that is invariant under translations of time [19].

In this process, the Fourier transform modulus is not exactly preserved for time series of finite length. Instead, we can observe a shift in the surrogate's Fourier transform modulus towards white noise [Fig. 3(b)]. This is because an estimate of the instantaneous measurement function will not be exact for time series of finite length, and consequently, small deviations arise that add uncorrelated random numbers to the data [20]. As a consequence, also the variance of the Fourier transform modulus frequency increases.

The usual problem of the method of surrogate data is that of a composite null hypothesis [21]. However, in our case, we account for non-Gaussianity by the usage of normally distributed random numbers in the process of generating the

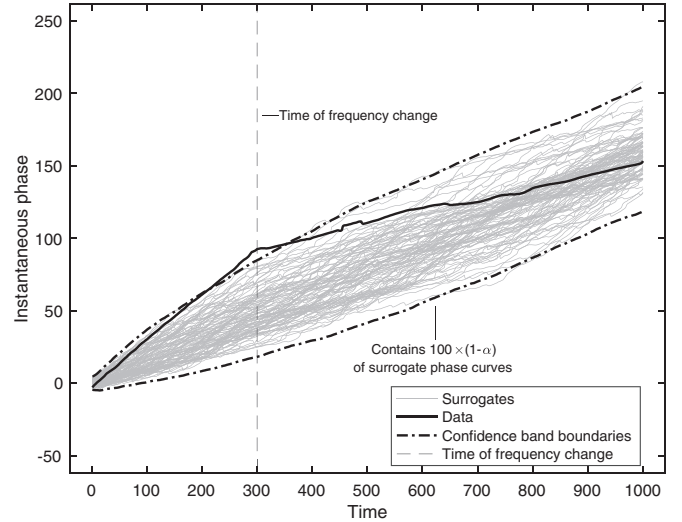


FIG. 4. Instantaneous phase of nonstationary example time series (black) from Fig. 2(a) and of 500 stationary surrogate time series (gray). Confidence bands for stationarity (dashed-dotted line) corresponding to a confidence level of $\alpha = 0.05$ are derived from the latter. Only a subset of surrogate phases is shown for better visibility.

AAFT surrogate time series. Nonlinearity can be detected, e.g., from the correlation structure of the bispectrum [22] or the presence of higher harmonics in the power spectrum. There were no signs for nonlinearity in the data we analyzed (Sec. II E). Therefore, the only remaining property in the null hypothesis is nonstationarity, which is also the property we want to test for.

C. Confidence bands are derived from surrogate phases

We test a time series for stationarity by comparing its instantaneous phase to that of a number of surrogate time series. This is done by deriving confidence bands from the latter.

A confidence band (Fig. 4) is constructed from an ensemble of instantaneous phases of surrogate time series as follows: We remove the mean from the ensemble and then calculate the $\beta/2 = 0.2$ and $1 - \beta/2$ quantiles at each time point. These two curves serve as preliminary lower and upper boundaries, which correspond to a pointwise confidence band. They neglect the fact that a trajectory that is extremal with respect to the ensemble only for a short period of time must be excluded from the confidence band as a whole. Therefore, these boundaries are too narrow and result in a size of the test larger than the significance level α . Size is the probability to incorrectly reject the null hypothesis if it is true, i.e., the probability of a type I error. We therefore adjust the confidence band by iteratively widening its boundaries: We scale lower and upper boundaries with a common factor γ until they contain $\geq 100 \times (1 - \alpha)\%$ of all surrogate phase curves, i.e., have the correct size. The null hypothesis of stationarity is rejected if the instantaneous phase of the test time series is outside of the confidence band confined by the lower and upper boundaries at some point in time. We also tested having individual scaling factors for the lower and upper boundaries. However, this did not yield a higher power of the test.




FIG. 5. Evaluation of the size and power of the test. For different frequency modulation strengths c_v [horizontal axis; Eq. (5)], we simulated $n = 2000$ realizations of a second-order autoregressive process and counted the rate of rejections of the null hypothesis (vertical axis) of the proposed test at a significance level of $\alpha = 0.1$. The test we propose is able to identify nearly 100% of nonstationary time series for $c_v \geq 1$ and does not lead to a fraction of false positive results larger than $\alpha = 0.1$ for $c_v = 0$.

The width of the confidence band increases with time. The ensemble of instantaneous phases has a maximum variability of 2π at $t = 0$. With every unwrapping at a phase jump of $\geq 2\pi$, differences in the instantaneous phase between individual surrogate time series are accumulating and lead to a widening of the ensemble.

D. Evaluation of size and power

To analyze the performance of the test we propose, we calculated the size and power of the test in a simulation study. Power is the ability of the test to correctly reject the null hypothesis if it is false, i.e., $1 - [\text{probability of type II error}]$. To estimate these two quantities, we simulated $n = 2000$ time series with different degrees of frequency modulation and therefore nonstationarity and counted the relative rate of rejections of the null hypothesis as a function of the frequency change (frequency modulation strength c_v ; Fig. 5).

As test data, we used realizations of an oscillating second-order autoregressive process [Eq. (1)] with a mixing time of $\tau = 20$ and 1100 time points, but discarded the first 100 time points, which corresponds to 5τ . This was done to make sure that correlation with the initial values decayed to a sufficiently small value. To introduce nonstationarity by modulating the frequency, we replaced ν in Eq. (2) with

$$\nu_0 \gamma(t), \quad \gamma(t) = \begin{cases} 1, & t \leq 500, \\ 2^{c_v}, & t > 500. \end{cases} \quad (5)$$

With a significance level of $\alpha = 0.1$, a base frequency of $\nu_0 = 1/20$, and 500 surrogate time series, we found that for stationary time series ($c_v = 0$, no frequency modulation), the rejection rate is almost 0.1, matching the significance level. The rejection rate increases with the frequency modulation

strength c_v , passes 75% at a frequency increase of $2^{0.68} - 1 \approx 60\%$, and reaches nearly 100% at $2^1 - 1 = 100\%$ change relative to the base frequency (Fig. 5).

We also tested the size and power of the test when varying the number of surrogate time series, mixing time τ , or length of the time series. Furthermore, we evaluated the impact of changing the frequency to smaller instead of larger values or using a chirp signal instead of the step-function modulation.

When varying the number of surrogate time series used to calculate the boundaries of the confidence band, we found that using more than 200 surrogate time series does not substantially increase the power of the test [Fig. 6(a)].

We also found that the power of the test decreases when we decrease the mixing time [Fig. 6(b)]. Low mixing times correspond to a fast decay of the autocorrelation function of the underlying process. In the limit case of $\tau = 0$, which represents white noise, all of the correlation structure including any possible oscillations and frequency modulation vanishes. The drop in power can therefore be attributed to the data itself, instead of being a shortcoming of the test we propose, and is inherent to all statistical methodology applied to the data.

Concerning the length of the time series, as expected, the power of the test in general benefits from longer time series [Fig. 6(c)]. At the same time, very short time series that span over only a few oscillations make the power decrease close to the level of significance regardless of the frequency modulation.

The test we propose also works well when the frequency is decreasing instead of increasing. As one would expect, we find a symmetric rejection rate with respect to positive versus negative frequency changes [Fig. 6(d)].

We observe that the power of the test is maximized when the two frequency regimes have equal length [Fig. 6(e)]. Making the length of the two regimes unequal reduces the power of the test, because it makes the frequency modulation less pronounced, and ultimately removes the frequency modulation from the data when one of the regimes has length zero. However, switching the frequency in the first half of the time series leads to an overall higher power of the test compared to switches in the second half. This is due to fluctuations accumulating in the instantaneous phase within the process of unwrapping, which increase the width of the confidence band, as discussed in Sec. II C.

Finally, we checked the power of the test when we use a chirp signal, i.e., replacing $\gamma(t)$ in Eq. (5) with

$$\gamma(t) = 1 + \frac{2^{c_v} t}{(\text{number of time points})}. \quad (6)$$

We find that the power of the test is similar to that when using a step-function frequency modulation [Eq. (5)]. However, in the latter case, the variance of frequencies is larger compared to the linear frequency increase with the same initial and final values in the chirp signal, which is why the power of the test increases less fast with c_v .

E. Application to experimental data

We applied our method to time series of HES/Her family protein expression in neural stem cells of zebrafish larvae that were captured using light sheet microscopy [Fig. 7]. These

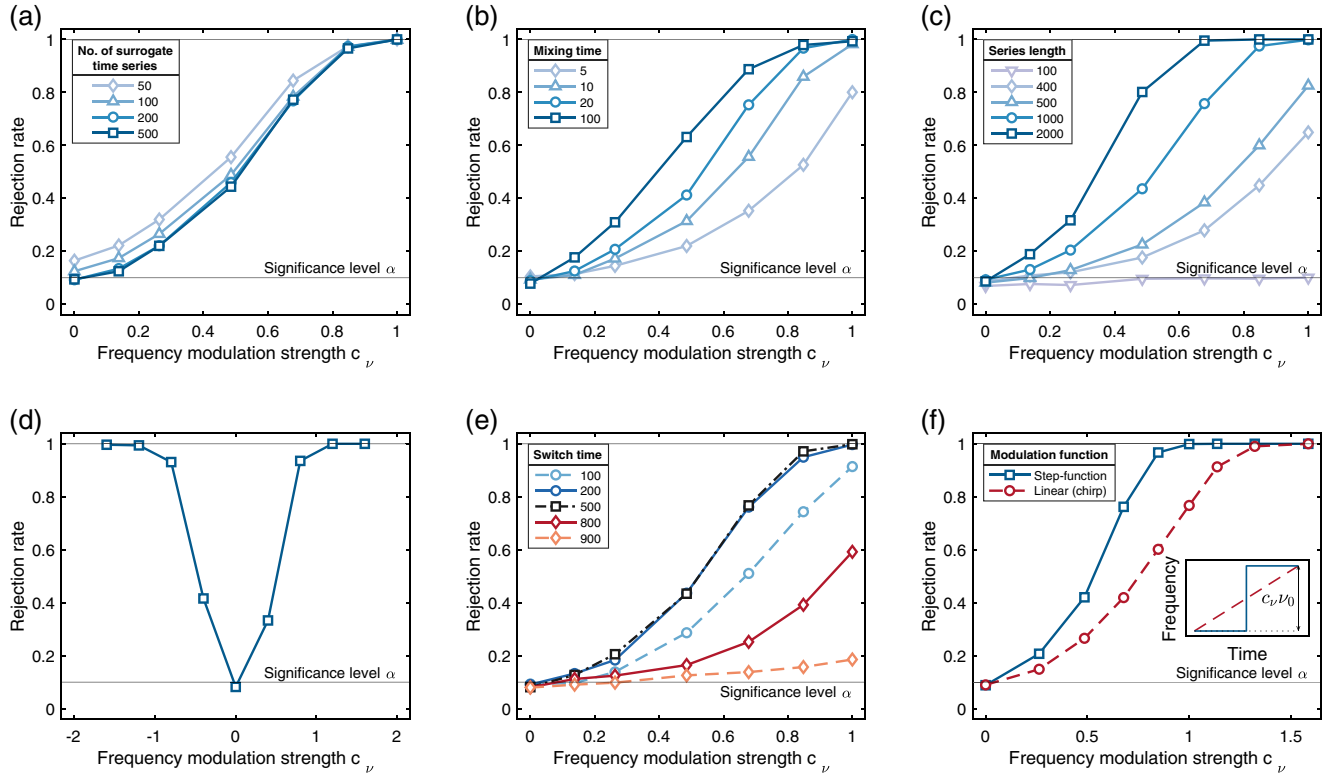


FIG. 6. We analyzed various characteristics of the simulated data and the test itself with respect to their influence on the power of the test (cf. Fig. 5) with $\nu_0 = 1/20$ and $\tau = 20$. (a) Power curves for different numbers of surrogate time series. (b) The mixing time τ in the simulated time series was tuned. (c) The influence of time-series length on the power of the test. (d) The effect of decreasing the frequency on the power of the test. (e) Power curves for different times of frequency modulation. (f) Compared with the step-function frequency modulation used in the other simulations, a chirp signal (linear modulation) with the same initial and final frequencies leads to an overall lower power. All time series used for the analyses had a length of 1000 time points if not indicated otherwise.




FIG. 7. We applied our method to the experimental data from Fig. 1(d). Confidence bands represent a confidence level of $\alpha = 0.05$. The test we propose rejects the null hypothesis of stationarity because the instantaneous phase of the experimental data exceeds the confidence band. Only a subset of the surrogate phases is shown for better visibility.

data feature rather short time series with only 86 time points and only a few oscillations due to the nature of the biological oscillation and imaging limitations *in vivo*. There were no signs for nonlinearity in the power spectrum (Sec. II B).

III. CONCLUSIONS

We propose a statistical test for nonstationary stochastic processes causing frequency modulation, especially in short time series. We evaluated its power and size and found that modulating the frequency by 80% and more relative to the base frequency leads to a correct identification of more than 97% of all nonstationary time series.

We tested our method with two functional forms of frequency modulation, a step function and a linear function. In a real-world scenario however, frequency modulation might come in a different form. We hypothesize that, for example, a fast oscillatory frequency modulation might not be detected by the test we propose, because deviations in the instantaneous phase from its linear appearance average out and the instantaneous phase does not exit the confidence band.

We demonstrated that the proposed test has power to detect frequency modulation in data that have only one frequency

at any given time. The superposition of signals with different frequencies but very similar amplitudes will decrease the performance of the method, but not lead to false positive results. Also, in data with extremely low mixing times or largely instationary sections, the power of our method is subpar. However, applying a statistical method to test for frequency modulation only makes sense if there are substantial portions of the data that portray differing frequencies, which excludes the extremal cases discussed in the simulation studies depicted in Figs. 6(e) and 6(f). Concerning unknown mixing and switch times, similar simulation studies can be conducted to determine the type II error, e.g., by estimating mixing times based on the autocorrelation functions.

The test we propose is tailored to detecting nonstationary processes in short time series that exhibit frequency modulation. Other types of nonstationarity might remain hidden to our method. Furthermore, there are no pure stationary processes in nature. Therefore, the null hypothesis will always be rejected for enough data [21]. However, our aim is detecting biologically relevant nonstationarity processes for which data are typically limited both in quantity and quality.

In summary, we propose a computationally cheap and easily interpretable test for frequency-modulated time series, which facilitates widespread usage and communication of results. It takes less than 1 s to analyze a series with 1000

time points on a state-of-the-art CPU (Intel Core i9-9880H). Additionally, there are next to no hyperparameters that have to be chosen dependent on the data except for the number of surrogate time series which shows a saturation behavior for its influence on the power of the test. The usage of surrogate data renders the test nonparametric, removing the need for a distribution assumption that can be violated, and allows the confidence bands being calculated in a way such that the test has the correct size by construction.

We believe that the method we propose provides a vivid meaning to nonstationarity in terms of frequency modulation and hope that this will further facilitate its application on real-world data.

An implementation of the described method for *Matlab* together with a demonstration script is available online [23].

ACKNOWLEDGMENTS

This study was supported by the German Research Foundation (DFG) under Germany's Excellence Strategy (CIBSS - EXC-2189 - Project ID 390939984) and SFB1381. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through Grant No. INST 35/1134-1 FUGG.

-
- [1] S. Codeluppi, L. E. Borm, A. Zeisel, G. La Manno, J. A. van Lunteren, C. I. Svensson, and S. Linnarsson, *Nat. Methods* **15**, 932 (2018).
 - [2] J. R. Moffitt, D. Bambah-Mukku, S. W. Eichhorn, E. Vaughn, K. Shekhar, J. D. Perez, N. D. Rubinstein, J. Hao, A. Regev, C. Dulac, and X. Zhuang, *Science* **362**, eaau5324 (2018).
 - [3] V. Svensson, R. Vento-Tormo, and S. A. Teichmann, *Nat. Protoc.* **13**, 599 (2018).
 - [4] T. Stuart and R. Satija, *Nat. Rev. Genet.* **20**, 257 (2019).
 - [5] R. Sueda, I. Imayoshi, Y. Harima, and R. Kageyama, *Genes Dev.* **33**, 511 (2019).
 - [6] B.-K. Liao, D. J. Jörg, and A. C. Oates, *Nat. Commun.* **7**, 11861 (2016).
 - [7] Y. Harima, Y. Takashima, Y. Ueda, T. Ohtsuka, and R. Kageyama, *Cell Rep.* **3**, 1 (2013).
 - [8] D. A. Dickey and W. A. Fuller, *J. Am. Stat. Assoc.* **74**, 427 (1979).
 - [9] D. Kwiatkowski, P. C. Phillips, P. Schmidt, and Y. Shin, *J. Econom.* **54**, 159 (1992).
 - [10] E. Zivot and D. W. K. Andrews, *J. Bus. Econ. Stat.* **20**, 25 (2002).
 - [11] P. C. B. Phillips and P. Perron, *Biometrika* **75**, 335 (1988).
 - [12] G. Elliott, T. J. Rothenberg, and J. H. Stock, *Econometrica* **64**, 813 (1996).
 - [13] J. Timmer, *Phys. Rev. E* **58**, 5153 (1998).
 - [14] S. D. Mohanty, in *Gravitational Waves: Third Edoardo Amaldi Conference*, edited by S. Meshkov AIP Conf. Proc. Vol. 523 (AIP, Melville, NY, 2000), pp. 479–480.
 - [15] J. Timmer, M. Lauk, W. Pflieger, and G. Deuschl, *Biol. Cybern.* **78**, 359 (1998).
 - [16] B. van der Pol, *J. Inst. Electr. Eng., Part 3* **93**, 153 (1946).
 - [17] D. Gabor, *J. Inst. Electr. Eng.* **93**, 429 (1946).
 - [18] L. Marple, *IEEE Trans. Signal Process.* **47**, 2600 (1999).
 - [19] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. Doyne Farmer, *Physica D* **58**, 77 (1992).
 - [20] T. Schreiber and A. Schmitz, *Physica D* **142**, 346 (2000).
 - [21] J. Timmer, *Phys. Rev. Lett.* **85**, 2647 (2000).
 - [22] U. Greb and M. G. Rusbridge, *Plasma Phys. Controlled Fusion* **30**, 537 (1988).
 - [23] <https://github.com/adrianhauber/FMtest>