

Tracking the Distance to Criticality in Systems with Unknown Noise

Brendan Harris^{*}

School of Physics, The University of Sydney, New South Wales 2006, Australia

Leonardo L. Gollo

*The Turner Institute for Brain and Mental Health, School of Psychological Sciences,
and Monash Biomedical Imaging, Monash University, Victoria 3168, Australia*

Ben D. Fulcher

School of Physics, The University of Sydney, New South Wales 2006, Australia

 (Received 25 October 2023; revised 17 April 2024; accepted 6 June 2024; published 8 August 2024)

Many real-world systems undergo abrupt changes in dynamics as they move across critical points, often with dramatic and irreversible consequences. Much existing theory on identifying the time-series signatures of nearby critical points, such as increased signal variance and slower timescales, is derived from analytically tractable systems, typically considering the case of fixed, low-amplitude noise. However, real-world systems are often corrupted by unknown levels of noise that can distort these temporal signatures. Here we aim to develop noise-robust indicators of the distance to criticality (DTC) for systems affected by dynamical noise in two cases: when the noise amplitude is either fixed or is unknown and variable across recordings. We present a highly comparative approach to this problem that compares the ability of over 7000 candidate time-series features to track the DTC in the vicinity of a supercritical Hopf bifurcation. Our method recapitulates existing theory in the fixed-noise case, highlighting conventional time-series features that accurately track the DTC. But in the variable-noise setting, where these conventional indicators perform poorly, we highlight new types of high-performing time-series features and show that their success is accomplished by capturing the shape of the invariant density (which depends on both the DTC and the noise amplitude) relative to the spread of fast fluctuations (which depends on the noise amplitude). We introduce a new high-performing time-series statistic, the rescaled autodensity (RAD), that combines these two algorithmic components. We then use RAD to provide new evidence that brain regions higher in the visual hierarchy are positioned closer to criticality, supporting existing hypotheses about patterns of brain organization that are not detected using conventional metrics of the DTC. Our results demonstrate how large-scale algorithmic comparison can yield theoretical insights that can motivate new theory and interpretable algorithms for solving important real-world problems.

DOI: [10.1103/PhysRevX.14.031021](https://doi.org/10.1103/PhysRevX.14.031021)

Subject Areas: Complex Systems,
Interdisciplinary Physics,
Nonlinear Dynamics

I. INTRODUCTION

A critical point, or bifurcation point, marks the value of some control parameter at which the dynamical properties of a system undergo a qualitative change, such as the appearance or disappearance of an attractor [1]. Many phase transitions occur at critical points described by bifurcations in macroscopic models of statistical ensembles

or complex systems [2]. In spin glasses, a net magnetization emerges at a critical value of the temperature order parameter [3], while in cosmology, the temperature controls the electroweak phase transition and baryon (matter-antimatter) asymmetry [4]. The dynamics of complex systems—from the brain to biological swarms—can also be understood in terms of their vicinity to a critical point, with near-critical systems possessing a range of functional advantages [5]. In the brain, for example, traversing critical points is thought to mediate changes in behavioral and cognitive functions such as memory storage and retrieval, as well as visual attention and processing [6–10].

Crossing a critical point can have dangerous consequences. For example, generator voltages in power distribution networks undergo bifurcation and transition from

^{*}Contact author: bhar9988@uni.sydney.edu.au

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

fixed-point stability to unsafe oscillation at a critical value of system load [11,12]. At so-called tipping points, where crossing a critical point leads to instability, the system may also change catastrophically as it rapidly transitions to a new state [13,14]. Catastrophes occur in simple physical systems, such as slipping events in wires under stress [15], as well as in models of sleep-wake transitions [16,17], transitions to epileptic brain states [18,19], and major climate events such as the desertification of North Africa [20]. Moreover, critical phenomena are studied across scientific domains, including physics [21,22], neuroscience [5,7], medicine [23,24], biology [25], and engineering [26]. The ubiquity of critical phenomena calls for a noise-robust method that can predict how close a system is to a critical point [27]: to give warning of epileptic seizure [18,19], anticipate power distribution failures [28], forecast imminent climate catastrophes [29], or to better control cognitive changes using brain stimulation [30].

Despite the ubiquity and diversity of critical systems, critical phenomena are remarkably linked by a common mathematical foundation: normal forms and bifurcation theory [13,14]. This theory categorizes most systems near a critical point into a well-studied, analytical taxonomy of bifurcations [1,31]. Moreover, a surprising number of systems exhibit relatively simple, codimension-one bifurcations [1,32], including saddle-node bifurcations found in the dynamics of neuronal spiking [33], pitchfork bifurcations in Bose-Einstein condensates [34], or Hopf bifurcations in laser cavities [35], financial markets [36], the auditory system [37], thalamocortical models [38], and models of interacting brain regions [39]. Each type of bifurcation is associated with a normal form—a dynamical equation that encapsulates the essential qualitative behavior near the critical point, regardless of the system [40]. Driven by the allure of solving a plethora of real-world challenges simultaneously, much recent work has studied this unifying mathematical structure of critical phenomena with the aim of identifying universal indicators of criticality. For normal forms and other systems that can be modeled analytically, this task is equivalent to estimating the control parameter that determines how close the system is to the bifurcation point. As such, a primary goal of studying normal forms has been to derive simple dynamical properties that diagnose how close a critical system is to a bifurcation point. Here, we term this quantity the distance to criticality (DTC) [41], which is determined by the control parameter in model systems.

Studies of normal forms and real-world critical systems have identified basic but universal dynamical characteristics that are informative of the DTC [42]: Near a critical point, dynamics are more variable and evolve on a slower timescale [25,43,44]. Both properties can be understood as a by-product of the flattening of the potential function of a system near the critical point. A flatter potential about a local minimum (fixed point) increases the return time when

a system is perturbed [45] and slows the timescale of fluctuations in a system driven by noise—a phenomenon known as critical slowing-down [43]. Because a flatter potential function also decreases the confinement of noise-related diffusion, the domain explored by a stochastic system will expand when the system is close to the critical point. Hence, the standard deviation of a time series increases near a bifurcation point, and has been shown to scale like a negative power of the control parameter when the strength of additive noise is close to zero [13,46,47]. Although properties of the potential function naturally motivate standard deviation and autocorrelation as indicators of criticality, many related features are used to estimate the DTC in a variety of fields [42,48] including time-series skewness of chlorophyll in lake ecosystems [49], increases in low-frequency fluctuations of voltage in electrical power systems [26], spatial autocorrelation of coupled dynamical systems [50], and other multivariate critical indicators [19,51,52]. Notably, qualitative signatures of criticality, such as maximized autocorrelation timescales, are universal across classes of bifurcations and forms of noise.

In deriving quantitative statistical indicators for the DTC, mathematical analyses typically assume the presence of a weak stochastic component relative to dominant deterministic dynamics [53]. Stochastic influences are typically classified as either measurement noise, which is applied to the final signal independent of the underlying dynamics of the system, or dynamical noise, which is incorporated into the equations of motion and constantly perturbs the deterministic dynamics [54,55]. Progress has been made in deriving scaling laws for critical indicators when the strength of dynamical noise is close to zero, such as the relationship between variance and the control parameter when approaching various bifurcations [13,14]. Unfortunately, even noise with a small but nonvanishing strength complicates the task of defining (let alone predicting) bifurcation points in stochastic systems [56,57]. One common approach, which we operate under in this work, is to designate the bifurcation point as the value of the control parameter that divides invariant densities with distinct qualitative properties (see Sec. III C). For example, if dynamical noise is incorporated into the normal form of a pitchfork bifurcation, then the distribution of states visited by the system converges over time for a fixed control parameter. The transition between a unimodal and bimodal invariant density corresponding to the appearance of a second attractor occurs at the same value of the control parameter as bifurcation in the deterministic terms alone [58]. This splitting of the invariant density (see Sec. III C) termed a phenomenological bifurcation [53] tends to precede other types of bifurcation points defined by changes in dynamical measures (known as dynamical bifurcations) [58] and gives a natural definition for the critical point that we use here.

Many physical systems, however, are subject to noise that is not small relative to the scale of their deterministic

dynamics. In real-world applications, the limitations of conventional indicators of nearby critical points are tied to the need for quantitative accuracy despite a limited understanding of the system, in particular, its stochastic components. Rather than looking for a qualitative increase or divergence in a statistical indicator of the DTC, which may be sufficient to diagnose a bifurcation after it has taken place, tasks such as estimating the amount of time (e.g., minutes or seconds) until a patient will have an epileptic seizure, or the amount of time (e.g., years) until an ecosystem will collapse, require precise calibration of the values of a critical indicator against a control parameter. In such cases, variations in the strength of noise destroy the precise relationships required for quantitatively estimating the DTC, such as by shifting the control parameter at which dynamical measures peak [58]. For example, moderate noise has been incorporated into models of critical transitions to neuron spiking [33], climate tipping [59], and the sleep-wake transition [17]. Although conventional critical indicators are analytically universal (close to the critical point, when noise is vanishingly small) and have been successfully applied to some real-world problems [18,20,42], they have yet to find widespread practical application in many noisy scenarios, such as predicting epileptic seizures [19,60]. Since properties such as standard deviation, autocorrelation, or skewness are highly sensitive to noise [47], the need for exact calibration is a point of failure for conventional indicators of criticality in systems where the strength of noise is unknown or variable. Dynamical noise—particularly when it has an unknown variance or distribution—is one of the major obstacles to applying universal theory and critical indicators derived for analytic systems to real-world scenarios involving finite, noisy time series.

To our knowledge, all conventional indicators of criticality are sensitive to noise. Some prior work has characterized how the strength of a noise process affects the time-series properties used to infer the control parameter, and therefore the DTC. Although scaling laws, which describe how the indicator varies with the control parameter, have been derived analytically for some indicators of criticality in low-noise cases, such a scaling may not hold under moderate or high levels of noise. For instance, Meunier and Verga [58] demonstrated for the pitchfork bifurcation that features such as the Lyapunov exponent and autocorrelation no longer peak at the phenomenological bifurcation point in the presence of additive noise. Instead, the peak of these time-series features (which marks a dynamical bifurcation point) is shifted along the control parameter to a degree that depends on the noise variance, thereby obscuring the critical point. Moreover, the splitting of the invariant density becomes less abrupt with increased noise, and other dynamical effects can occur after bifurcation of the underlying deterministic component [58,61]. Even a recent measure formulated to reliably signal the

critical point across classes of bifurcations is sensitive to dynamical noise [62].

While some studies have examined the behavior of critical indicators in systems with a given fixed noise level, prior work has yet to address real-world settings in which the noise amplitude is unknown and may vary (across different recordings of the system, or over instances of otherwise similar systems). For example, variable dynamical noise may arise when studying (i) the ecological dynamics of species, with noise varying between sites with different climates or weather patterns, (ii) the response of brain areas to random stimulation, where the strength of the stimulation is metered by anatomical differences such as skull thickness [63], (iii) the sleep-wake or epilepsy transition, in which dynamical noise may be influenced by a subject's individual physiology or environment [17,64], (iv) climate systems, in which the influence of sources of variability such as human activity change over time [59,65], or (v) components of other systems that may be influenced by different external stochastic drives or complex internal heterogeneity [9,66]. This problem is particularly challenging because knowledge of the noise process, especially its variance, is required to properly calibrate time-series features for meaningful prediction of the DTC. Additionally, established analytical tools that have been used to make progress in idealized cases cannot be applied straightforwardly to the variable-noise setting because their precise relationship to the control parameter depends on the noise amplitude. For example, two prototypical time-series indicators of the DTC—standard deviation and autocorrelation—both depend strongly on the strength of noise in the system: standard deviation increases with noise level, while autocorrelation decreases.

Given the ubiquity of systems corrupted by an uncertain amount of noise, it is of great importance to formulate robust indicators of the distance of such systems to nearby critical points. In this work, we introduce a data-driven methodology to tackle this problem, in which we simulate a noisy bifurcating system close to a critical point and then search across a large candidate library of time-series features, scoring each on how well it captures the DTC of a simulated system (Sec. II C). The most comprehensive collection of time-series features to date is the highly comparative time-series analysis library *hctsa*, which contains over 7000 diverse features [67–69]. It contains measures of the distribution of time-series values, self-correlation properties, measures of predictability and complexity, methods relating to self-similarity and recurrence properties derived from the literature on physics and nonlinear time-series analysis, among many others. Here, we address the challenge of inferring the DTC using the data-driven approach of identifying the most informative time-series features for a given criticality problem by searching across the *hctsa* feature set. To study how each time-series feature is influenced by dynamical

noise, we consider our model system in two scenarios (Sec. II C): (i) the fixed-noise case in which the noise has a fixed strength and (ii) the variable-noise case in which the noise amplitude varies across repeated measurements of the system. While ensemble methods that use combinations of time-series features to optimize a performance metric have been employed in many prior applications of feature libraries such as *hctsa* [67,70], we aim in this work to demonstrate an alternative approach: that a large library of time-series analysis algorithms can form the foundation for building novel time-series theory and interpretable algorithms for analyzing real-world data. We find the most noise-robust features by identifying those that covary most strongly with the control parameter in the case of variable noise, then use close inspection to explain how these high-performing features combine the characteristics of noise-driven fast fluctuations with measurements taken from the stationary distribution (see Sec. III C). This combination enables a precise estimate of the shape of the potential, and hence an ability to track the DTC, despite the confound of uncertain noise amplitude.

The paper is structured as follows. In Sec. II, we introduce the supercritical Hopf normal form that we use as a model critical system (in Sec. II A) and describe how we sample noise and control parameters to generate a dataset of simulated time series (in Sec. II B). We detail our feature-extraction procedure and our scoring method for identifying high-performing features in the fixed-noise and variable-noise cases in Sec. II C. Next, in Sec. III, we reveal how well features across the *hctsa* library perform in the fixed-noise (Sec. III A) and variable-noise (Sec. III B) cases. After identifying the top-performing *hctsa* features in the variable-noise case and examining them in detail (Sec. III C), in Sec. III D we summarize the theoretical insight we gained from our data-driven approach and encapsulate our findings in a new time-series feature, which we call rescaled autodensity (RAD), that can robustly track the DTC in near-critical systems with unknown noise amplitude. In Sec. III E, we use RAD to examine mouse electrophysiology data sourced from the Allen Neuropixels Visual Behavior dataset. Our investigation tests a structure-function hypothesis positing that regions higher in the cortical hierarchy exhibit longer timescales [71–75] due to their closer proximity to criticality [7,76]. Our findings demonstrate the ability of RAD to accurately track the anatomical hierarchy of visual cortical regions, outperforming conventional indicators of the DTC. Finally, in Sec. IV we provide concluding remarks, discuss the implications of our findings for other noisy critical phenomena, and outline potential directions for future research.

II. METHODS

Here we outline our data-driven approach to finding useful statistical indicators of the DTC in the presence of

dynamical noise in two cases: the fixed-noise case and the variable-noise case. Our approach is shown schematically in Fig. 1 and can be summarized in two main steps: (i) simulate a time-series dataset from the normal form of a supercritical Hopf bifurcation [see Eq. (1)], where each time series is generated using a specific control parameter (varying over a range up to the bifurcation point) and noise variance (either fixed or varying over a specified range), and (ii) extract candidate features from each time series using the *hctsa* time-series feature library [67,68] and score each feature on how well it tracks the control parameter μ across the dataset. In the fixed-noise case, the noise process has a constant strength; i.e., η is fixed for all time series in the dataset. By contrast, in the variable-noise case the noise amplitude can take a different value for any given time series; this setting models real-world situations in which the strength of noise might vary across different recordings of a system or across different instances of similar systems.

A. The model system

We first describe the model system used for generating time series at various distances, in the control parameter, from a critical point. We choose the radial component of the normal form for a supercritical Hopf bifurcation with dynamical noise, as it applies to a broad range of real-world systems, including the wake-sleep transition [17], seizure dynamics [64], auditory hair cells [77,78], financial markets [36], and many others [6,7,37,79,80]. Our system derives from the general normal form of a supercritical Hopf bifurcation [81], in which a stable fixed point bifurcates to a stable limit cycle and an unstable fixed point. Since noise often acts radially with a strength that is independent of the system’s phase, we choose to use the simplified radial component of the Hopf normal form given by

$$dx = (\mu x - x^3)dt + \eta dW, \quad x \geq 0, \quad (1)$$

where W is a Wiener process. Notably, this resembles the normal form of pitchfork bifurcation with a reflecting boundary at the origin [6], permitting our results to generalize via a simple transformation to systems that exhibit pitchfork-type bifurcations (such as spin-glass systems [3] or models of epigenesis [82]). As shown in the bifurcation diagram in Fig. 1(a), a bifurcation occurs at the critical value $\mu = 0$: For $\mu < 0$ there is a single stable fixed point at $x = 0$, whereas for $\mu > 0$ the origin is unstable. The corresponding potential function of this system detailed in Sec. III C is unimodal for $\mu < 0$, flattens as μ is increased to 0, and is bimodal for $\mu > 0$.

Here we consider the more complex problem of approaching the critical point from $\mu < 0$ corresponding to the regime with a single, stable fixed point where the system hovers near the origin [shaded in Fig. 1(a)]. We focus our study on this regime for three reasons. First,

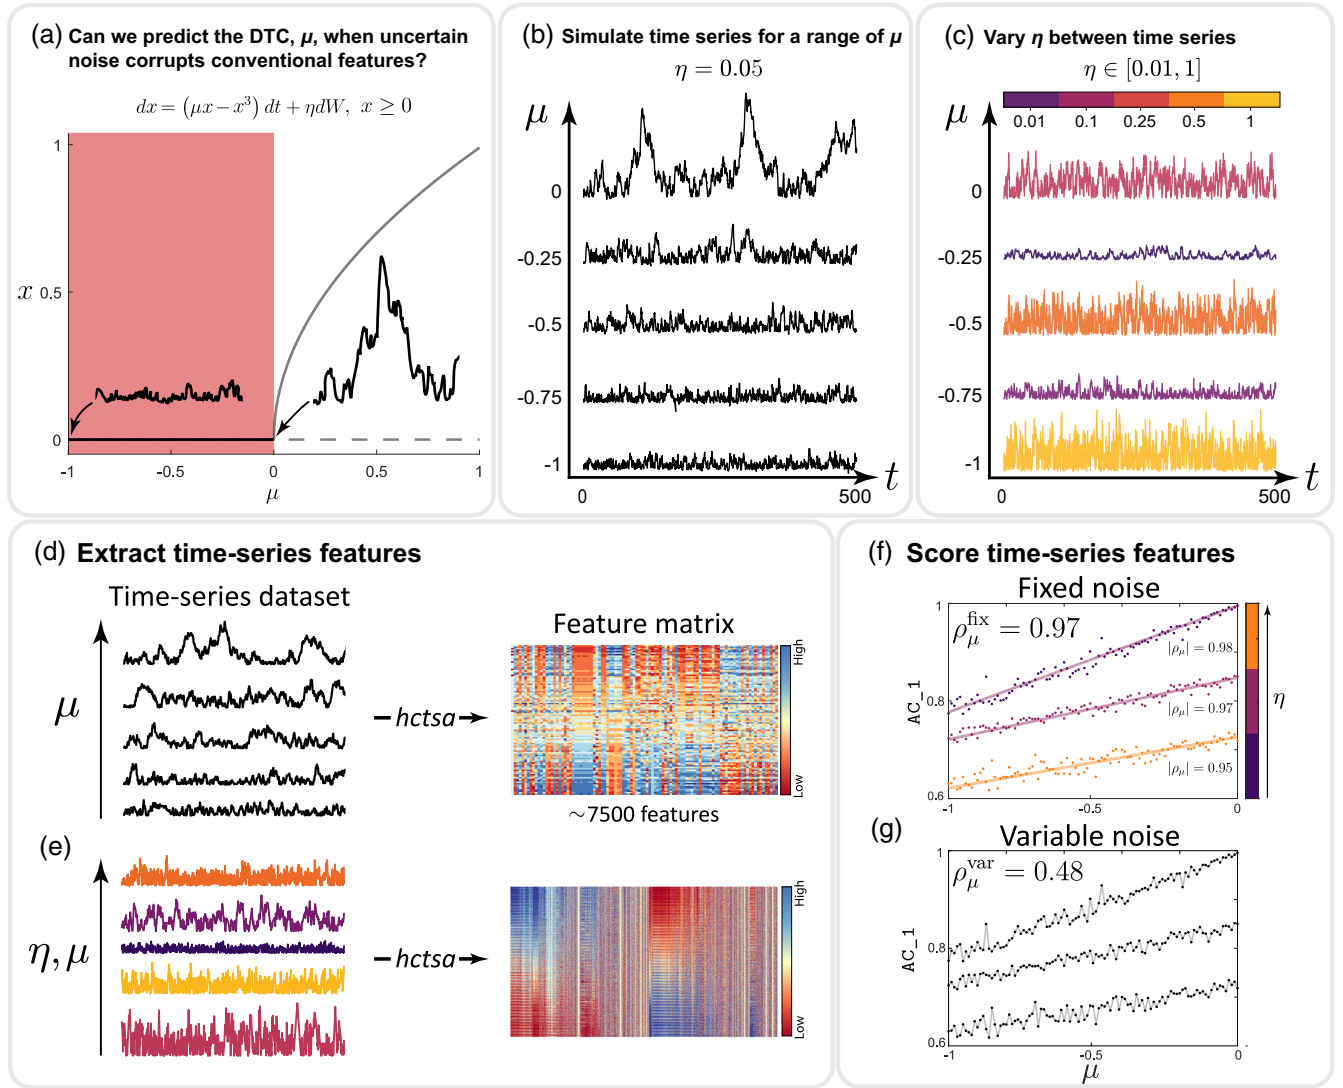


FIG. 1. We take a data-driven approach to identifying time-series features that accurately track the DTC of a noisy, near-critical system across a range of noise levels. (a) We investigate a model system close to the critical point: the radial component of the normal form for a supercritical Hopf bifurcation with $\mu < 0$ given by Eq. (1). A bifurcation diagram shows the radius of equilibrium points x against the control parameter μ ; when $\mu < 0$ (shaded red), there is a stable fixed point (solid black) at the origin. The model incorporates additive white Gaussian noise given by dW with a noise amplitude of η . We simulate time series across $-1 \leq \mu \leq 0$ and $0 < \eta \leq 1$. Snippets of representative time series are annotated. (b) Representative time series of 5000 samples for the fixed-noise problem, for which the noise level is fixed (shown here for $\eta = 0.01$) at selected values of $\mu = -1, -0.75, \dots, 0$. As μ approaches zero, it is visually clear that signal variance increases and fluctuations are slower. (c) Sample time series for the variable-noise problem: A different value of $0.01 \leq \eta \leq 1$ is chosen for each time series, as indicated by the coloration. Relative to the fixed-noise case, the variable-noise amplitude obscures the underlying variation in the control parameter μ . (d),(e) To find time-series properties that are sensitive to μ in both the fixed-noise and variable-noise settings, we compare 7492 candidate time-series features using *hctsa* [68]. (d) The result of this feature extraction for a fixed value of η is depicted as a time series \times feature matrix, with rows ordered by μ . (e) The feature matrix for our full dataset, with variation in both μ and η , is depicted in the right panel, with rows ordered by η . (f),(g) Finally, we score each feature for its performance under the two conditions as ρ_{μ}^{fix} and ρ_{μ}^{var} . In the fixed-noise setting (f), ρ_{μ}^{fix} is computed by averaging the magnitudes of the correlation of a feature with μ for each separate value of $\eta = 0.01, 0.02, \dots, 1$ (colored lines). In the variable-noise case (g), ρ_{μ}^{var} is an overall correlation to μ after η labels are discarded (indicated in gray). An example is shown here for the autocorrelation at lag 1 (labeled AC_1), which is strongly correlated to μ for a given noise-level η (f) but weakly correlated only in the variable-noise case (g).

during events such as power-system failure [11,12] and epileptic seizures [64] the fixed-point regime corresponds to safety, with the presence of low-power fluctuations, in contrast to the dangerous high-power oscillations that occur

for $\mu > 0$. Second, the DTC is much more straightforward to estimate for $\mu > 0$: Features that measure simple properties of the distribution, such as the mean or median, are insensitive to η when the noise amplitude is not large

compared to the equilibrium radius [given by $\sqrt{\mu}$ for Eq. (1)]. Third, when the noise strength is large enough to mask changes in the distribution caused by varying $\mu > 0$, the system is already sufficiently close to the bifurcation point for the disappearance of the unstable branch to be insignificant.

B. Time-series simulation

To evaluate the performance of individual time-series features at tracking μ in the presence of a stochastic component with strength η , we simulated Eq. (1) in the range $-1 \leq \mu \leq 0$. We studied both the fixed-noise and variable-noise scenarios with a combined time-series dataset generated by varying parameters across ranges $-1 \leq \mu \leq 0$ and $0 < \eta \leq 1$. In total, we simulated 10100 time series \mathbf{x} using all combinations of 101 equally spaced values for $\mu = -1, -0.99, \dots, 0$, and 100 equally spaced values of $\eta = 0.01, 0.02, \dots, 1$. Time series were simulated using the Euler-Maruyama method [83] over 1000 s with a time step of 10^{-3} s. To avoid the effects of transient dynamics (which are sensitive to the initial condition), we set an initial radius of $x_0 = 0$ coinciding with the stable fixed point of the system. Moreover, we discarded the first 500 s of the integration period corresponding to an interval far wider than the longest timescale in any of our simulations (estimated using the first zero crossing of the autocorrelation function). Finally, we down-sampled the remaining 500 s to a sampling period of $\Delta t = 0.1$ s, yielding 5000-sample time series that were analyzed in the remainder of this work. From this combined dataset, we studied the fixed-noise case by searching for features that varied with μ when η was fixed, as depicted in Fig. 1(b). In the variable-noise case depicted in Fig. 1(c), we searched for time-series features that varied with μ in a way that was consistent over confounding variation of η . Finally, note that our data-driven approach was focused on relatively short time series containing 5000 samples, in contrast to other studies that evaluated critical indicators on simulated data [62]. Representative examples of simulated time series are in Figs. 1(b) and 1(c).

C. Feature scoring

We next aimed to determine which types of time-series features extracted from the noisy time series simulated above could accurately track variations in the known parameter μ . We achieved this in a data-driven way using a comprehensive collection of 7873 candidate time-series features from the highly comparative time-series analysis toolbox *hctsa* (version 0.98) [67,68]. Each time-series feature $f: \mathbb{R}^N \rightarrow \mathbb{R}$ maps an input time series (of $N = 5000$ samples here) to a single, real-valued summary statistic. The *hctsa* feature set contains methods developed across the interdisciplinary time-series analysis literature, including measures of outliers, periodicity, stationarity,

predictability, self-affine scaling, and many others. We extracted all *hctsa* features from each simulated time series in the dataset described above. After feature extraction, we removed 381 features from our analysis that were not well behaved across all time series (e.g., produced not-a-number values or constant outputs across the dataset), leaving 7492 good-quality features. The result of feature extraction across our full time-series dataset is visualized as a time series \times feature matrix in Fig. 1(e).

Our next goal was to assess the ability of each time-series feature to track the underlying variation of the control parameter μ across the time-series dataset described in Sec. II B. We scored each feature using a Spearman correlation coefficient ρ between its outputs across the dataset and the corresponding values of μ ; features with high $|\rho|$ strongly monotonically track the variation in μ . As depicted in Figs. 1(f) and 1(g), we considered the fixed-noise and variable-noise cases separately.

1. Fixed-noise case

In the fixed-noise case, where the noise amplitude is constant across recordings, we analyzed subsets of time series with the same value of η . As depicted in Fig. 1(f), for each set of 100 time series with a given value of η , we calculated the Spearman correlation of each feature with μ (over $\mu = -1, -0.99, \dots, 0$). We then computed the fixed-noise score of a given feature ρ_μ^{fix} as the mean magnitude of these 100 Spearman correlation coefficients calculated independently for each noise level. That is, ρ_μ^{fix} measures the average performance of a time-series feature across many settings that each have a unique fixed noise level.

2. Variable-noise case

In the variable-noise case, where the noise amplitude varied between recordings, we disregarded the η labels of each time series (treating them as unknown), as depicted in Fig. 1(g). We then computed the variable-noise score of a given feature ρ_μ^{var} as the Spearman correlation between the 10 100 feature values (over all time series) and the corresponding values of μ . In contrast to Fig. 1(f), all points in Fig. 1(g) are uncolored, reflecting how the variable-noise score is a correlation across all data points, regardless of the noise strength. That is, ρ_μ^{var} measures the performance in a single setting where the noise amplitude varies across time series. Unlike the fixed-noise score ρ_μ^{fix} (which takes absolute values and is therefore non-negative by construction), the variable-noise score ρ_μ^{var} can be positive or negative, indicating whether the feature increases or decreases with the proximity to criticality.

III. RESULTS

By systematically exploring a supercritical Hopf bifurcation, our main aim was to investigate which types of

time-series properties (from a diverse library of over 7000 candidates in *hctsa* [68]) accurately track the DTC given by $|\mu|$ in the precritical regime of $\mu < 0$. As detailed in Sec. II C above, features were scored based on their correlation to μ using either ρ_μ^{fix} in the fixed-noise setting (Sec. III A), and, separately, ρ_μ^{var} in the variable-noise setting (Sec. III B). In the variable-noise case, we found surprising types of features that strongly varied with the DTC. By examining the algorithmic elements that are common among these features and summarizing the theoretical insight they provide (Sec. III C), we developed a new time-series feature for tracking the DTC, which we named the rescaled autodensity (see Sec. III D).

A. The fixed-noise case

We first investigated the fixed-noise case, where individual features were scored according to their correlation with μ for a fixed noise level η . We quantified this correlation as ρ_μ^{fix} (an average of absolute Spearman correlation coefficients across each of the 100 noise levels η ; cf. Sec. II C). We hypothesized that features that are well known to track the vicinity of a system to a critical point—including measures of signal variance and

autocorrelation—would receive high ρ_μ^{fix} scores. We also sought to investigate whether any new types of features showed strong performance.

The distribution of ρ_μ^{fix} values across all 7492 time-series features is shown as a histogram in Fig. 2(a). We found many time-series features with strong performance on this fixed-noise task, e.g., 1348 of the 7492 tested features have $\rho_\mu^{\text{fix}} > 0.8$, and some features were scored as high as $\rho_\mu^{\text{fix}} = 0.98$. To more closely investigate the highest-performing time-series features, we focused on the 100 features with $\rho_\mu^{\text{fix}} > 0.97$. We then sought to isolate groups of similarly behaving features from within this set of 100 features by computing pairwise Spearman correlation coefficients between all pairs of features (calculated using all time series), and reordering them using hierarchical linkage clustering. The results shown in Fig. 2(b) reveal two distinct groups of high-performing features: one group measuring autocorrelation properties (annotated with an orange square), and another (smaller) group measuring properties of the distribution of time-series values, including mean and variance (annotated with a blue square).

The first cluster contains features that are sensitive to autocorrelation; all features in this cluster are highly correlated to lag-1 autocorrelation (*hctsa* feature name

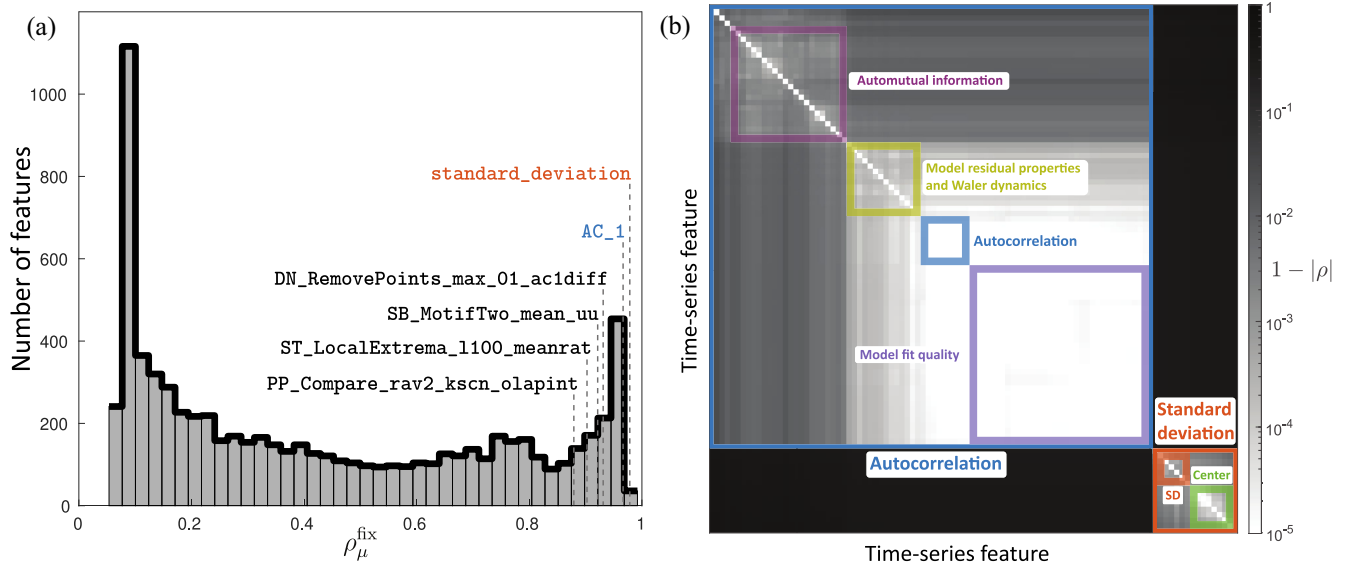


FIG. 2. For a near-critical system with a fixed noise level, we recovered high-performing conventional statistical indicators of the distance to criticality by comparing the behavior of over 7000-candidate time-series features. (a) A histogram of fixed-noise feature scores ρ_μ^{fix} across 7492 time-series features (from the *hctsa* feature library). Selected high-performing features are annotated including standard deviation ($\rho_\mu^{\text{fix}} = 0.98$, labeled as `standard_deviation`, blue) and lag-1 autocorrelation ($\rho_\mu^{\text{fix}} = 0.97$, `AC_1`, red). Names of the top features for the variable-noise case are also annotated for comparison (cf. Fig. 3). (b) The top 100 features with $\rho_\mu^{\text{fix}} \geq 0.97$ are plotted as a pairwise correlation matrix. The brightness of each element corresponds to the similarity between each pair of these features, using the metric $1 - |\rho|$, where ρ is the Spearman correlation coefficient. Rows and columns have been reordered to place similar features close to one another using average hierarchical linkage clustering (on Spearman correlation distances), revealing two clusters of features with common behavior, which have been annotated using transparent colored squares. As labeled, the first cluster contains features measuring properties of signal autocorrelation (orange), while the second cluster contains features measuring properties related to the signal variance (blue). Features annotated in (a) are detailed in Table I, and for (b) a sorted list of features is in Supplemental Material Table S2 [84].

AC_1), which has $\rho_\mu^{\text{fix}} = 0.97$. Disregarding the confounds of noise amplitude and the time step, lag-1 autocorrelation and other features that measure timescales of a system are sensitive to the critical slowing-down that occurs close to the critical point [20]. As well as lag-1 autocorrelation, our data-driven analysis also revealed a range of conceptually related features that can also effectively capture the same variation in the self-correlation timescale, including measures of automutual information [85], properties of fitted autoregressive time-series model residuals, and others (see Supplemental Material Table S2 [84] for a list of the top 100 fixed-noise features clustered by similarity).

In the second cluster of features, which are related to the distribution of time-series values, the standard deviation (*hctsa* feature name `standard_deviation`) displays very strong performance with $\rho_\mu^{\text{fix}} = 0.98$. This is consistent with expectation; the potential function described in Sec. III C flattens when approaching a critical point, leading to higher-variance fluctuations for a given noise level η . Moreover, since the potential for our model system has a reflecting boundary at $x = 0$ [see Eq. (2)], we also found that measures of central tendency such as the mean, median, and root-mean-square, are highly correlated to the spread of the distribution (as captured by the standard deviation).

Our results have thus flagged two types of features that vary with the DTC, recapitulating prior literature which has focused on critical slowing-down (captured by measures of time-series autocorrelation [20], including the simple and

effective indicator, lag-1 autocorrelation [13,14,42,43,48]) and fluctuations of increased variance (captured by the spread of the distribution of time-series values) [86] near a critical point. A complete list of *hctsa* features along with their fixed-noise scores ρ_μ^{fix} is in Supplemental Material Table S1 [84].

B. The variable-noise case

Having established the ability of our data-driven approach to recapitulate a theoretical literature on time-series features for tracking the DTC in a system with fixed noise, we next investigated the more difficult problem of finding features that track the DTC in a variable-noise setting. We are unaware of prior work that has examined this relevant real-world scenario with a comprehensive evaluation of statistical properties of time series that are sensitive to μ but insensitive to η . The lack of an *a priori* understanding of how to construct a noise-level-robust indicator of the DTC from finite time series, and the difficulty (or intractability) of a direct analytical route to tackling the problem, makes it an ideal setting for our data-driven approach.

We computed the variable-noise performance score ρ_μ^{var} for all candidate time-series features in *hctsa* (full results are in Supplemental Material Table S1 [84]). The distribution of $|\rho_\mu^{\text{var}}|$ across all 7492 features is plotted in Fig. 3(a), including annotated scores of two selected features (the standard deviation and lag-1 autocorrelation).

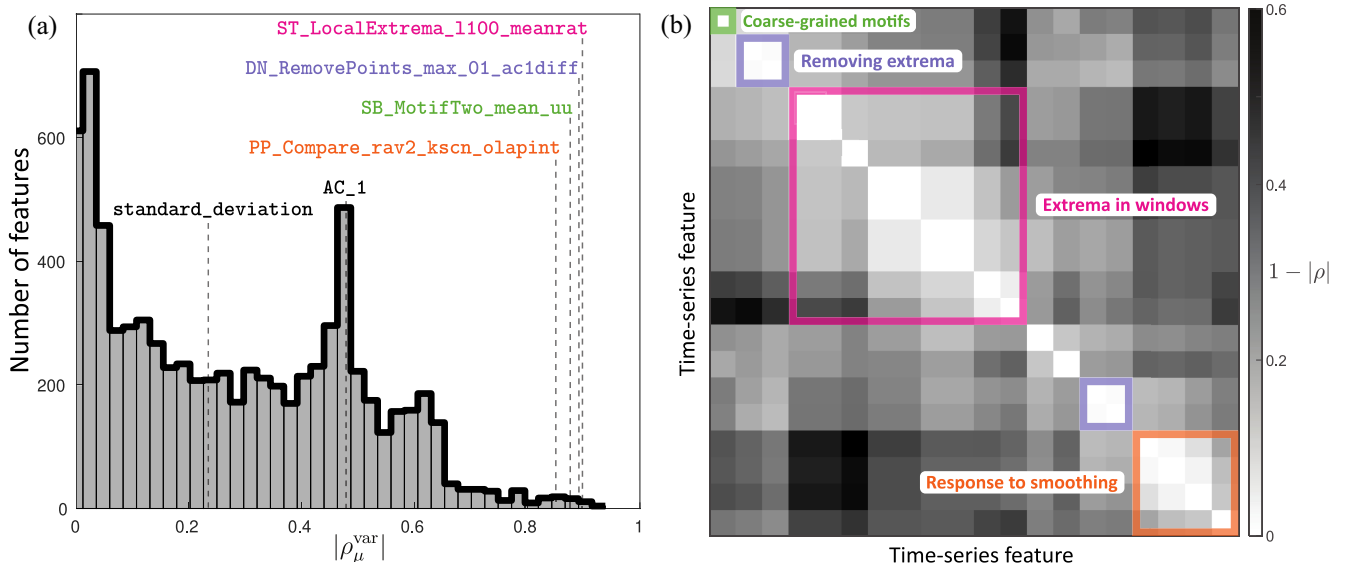


FIG. 3. In the variable-noise setting, we identified a set of high-performing time-series features that measured properties of the distribution of time-embedded values. (a) A histogram of the variable-noise scores $|\rho_\mu^{\text{var}}|$ is shown across all 7492 *hctsa* time-series features. Conventional indicators of criticality perform poorly, including standard deviation ($\rho_\mu^{\text{var}} = 0.23$, annotated as `standard_deviation`) and lag-1 autocorrelation ($\rho_\mu^{\text{var}} = 0.48$, annotated as `AC_1`). (b) The top 20 features in the variable-noise case ($|\rho_\mu^{\text{var}}| > 0.86$) are plotted in a pairwise correlation matrix, revealing many diverse clusters of features. Families of conceptually similar features are annotated, and the top features from outlined clusters were selected for closer study. Features annotated in (a) are detailed in Table I, and a sorted list of the features in (b) is in Supplemental Material Table S3 [84].

Reflecting the increased difficulty of the variable-noise setting, $|\rho_\mu^{\text{var}}|$ values are lower on average than ρ_μ^{fix} . We nevertheless observed a tail of high-scoring features; e.g., 49 time-series features have $|\rho_\mu^{\text{var}}| > 0.8$. As expected, the top-performing features in the fixed-noise case, including lag-1 autocorrelation (`AC_1`) and standard deviation (`standard_deviation`) are sensitive to both μ and η , and thus do not accurately track variation in μ in the presence of confounding variation in η . For example, for `standard_deviation`, $\rho_\mu^{\text{fix}} = 0.98$ (in the fixed-noise setting) drops to $\rho_\mu^{\text{var}} = 0.23$ (in the variable-noise setting), while for `AC_1`, $\rho_\mu^{\text{fix}} = 0.97$ drops to $\rho_\mu^{\text{var}} = 0.48$. To understand this behavior, we plotted the dependence of these two features as a function of μ for five selected noise levels in Figs. 4(b) and 4(a). We saw that these features varied monotonically with μ for any given η value (underlying their strong fixed-noise performance), but they were highly sensitive to variation in η , making them unreliable indicators of the DTC in the variable-noise setting.

We next aimed to better understand the high-performing features that make up the right tail of Fig. 3(a). Focusing on the top 20 features, with $\rho_\mu^{\text{var}} > 0.86$, we plotted their pairwise distance matrix (on Spearman correlation distances, $1 - |\rho|$) in Fig. 3(b) (see Supplemental Material Table S3 [84] for details on each cluster). One large family of highly correlated features examines how properties of extreme values within small windows (of approximately 100 samples) are distributed across the time series, such as the average ratio between the maximum and minimum value in each window [`ST_LocalExtrema_1100_meanrat` shown in pink in Fig. 2(a)]. Other high-performing features include those that measure the change in autocorrelation after a proportion of time-series extrema are removed (such as `DN_RemovePoints_max_01_ac1diff`, purple), the occurrence of simple motifs in a symbolization of the time series (i.e., transforming a sequence of real values to a symbolic string; `SB_MotifTwo_mean_uu`, green), or the change in the distribution of time-series values after smoothing the time series using a moving average (`PP_Compare_rav2_kscn_olapint`, orange).

C. Algorithmic steps underlying noise-robust features

In the previous section, we isolated a set of time-series features that can track the control parameter μ , while being minimally affected by changes in the noise level η , with a surprisingly strong correlation ($\rho_\mu^{\text{var}} > 0.86$). In this variable-noise setting, conventional indicators of the DTC, like standard deviation and lag-1 autocorrelation, perform poorly. In this section, we aim to understand why these features perform so well. We are able to explain this by inspecting the algorithms underlying four of the highest-performing features selected to represent the main clusters of high-performing feature behavior [in Fig. 3(b)]. These

four features, which are named and briefly described in Table I, are annotated in Figs. 2(a) and 3(a). Scatter plots of these top features with the DTC μ for different noise levels η [plotted in Figs. 4(c)–4(f)], as well as the joint distribution of *hctsa* features over the fixed-noise and variable-noise scores [plotted in Supplemental Material Fig. S1(a) [84]] indicate their high performance in both the fixed-noise and variable-noise cases. Unlike the standard deviation [Fig. 4(a)] and lag-1 autocorrelation [Fig. 4(b)], these high-performing features vary strongly with μ in a similar way across noise levels η , demonstrating their robustness to noise.

Despite appearing to be distinct features, close inspection of the algorithmic steps underlying each top feature revealed two key shared algorithmic components: (i) They involve estimating a statistic derived from the distribution of time-series values, such as a proportion of points within an interval of values, and (ii) they compare this statistic to the standard deviation of the incrementally differenced time series (a quantity that is closely related to the lag-1 autocorrelation). In this section, we aim to summarize how these two algorithmic steps, which appear critical to the top-performing features, allow them to act as noise-robust estimators of the DTC. As with conventional metrics of criticality, such as standard deviation and autocorrelation, we notice that both algorithmic steps are closely related to a system's potential function. By expressing these two properties in terms of the potential, we are able to explain how these features can successfully infer the DTC despite uncertain noise amplitude. We first derive the relationship between the potential and the two algorithmic steps: (i) the distribution of values and (ii) the spread of differences.

1. The shape of the potential $V(x)$ depends on μ

The ability of some features to accurately track the DTC in the presence of uncertain noise amplitude can be understood in terms of the potential function of a system. The potential function of a stochastic system describes the effect of the deterministic components on the system dynamics: The gradient of the potential at a given point determines the rate of change due to the deterministic terms at that point [1]. In this formalism, the evolution of a system can be viewed intuitively as the trajectory of a heavily damped particle following the gradient of the potential toward a local minimum [1] in addition to any diffusive (or stochastic) drives. The potential formulation provides a useful way of understanding critical slowing-down and other conventional metrics of the DTC [42]. The potential function for the model system studied here is obtained by integrating Eq. (1), giving

$$V(x; \mu) = -\frac{\mu x^2}{2} + \frac{x^4}{4}, \quad x \geq 0. \quad (2)$$

As depicted in Figs. 5(a)–5(d), as μ increases toward 0, $V(x)$ becomes shallower, reducing the restorative force

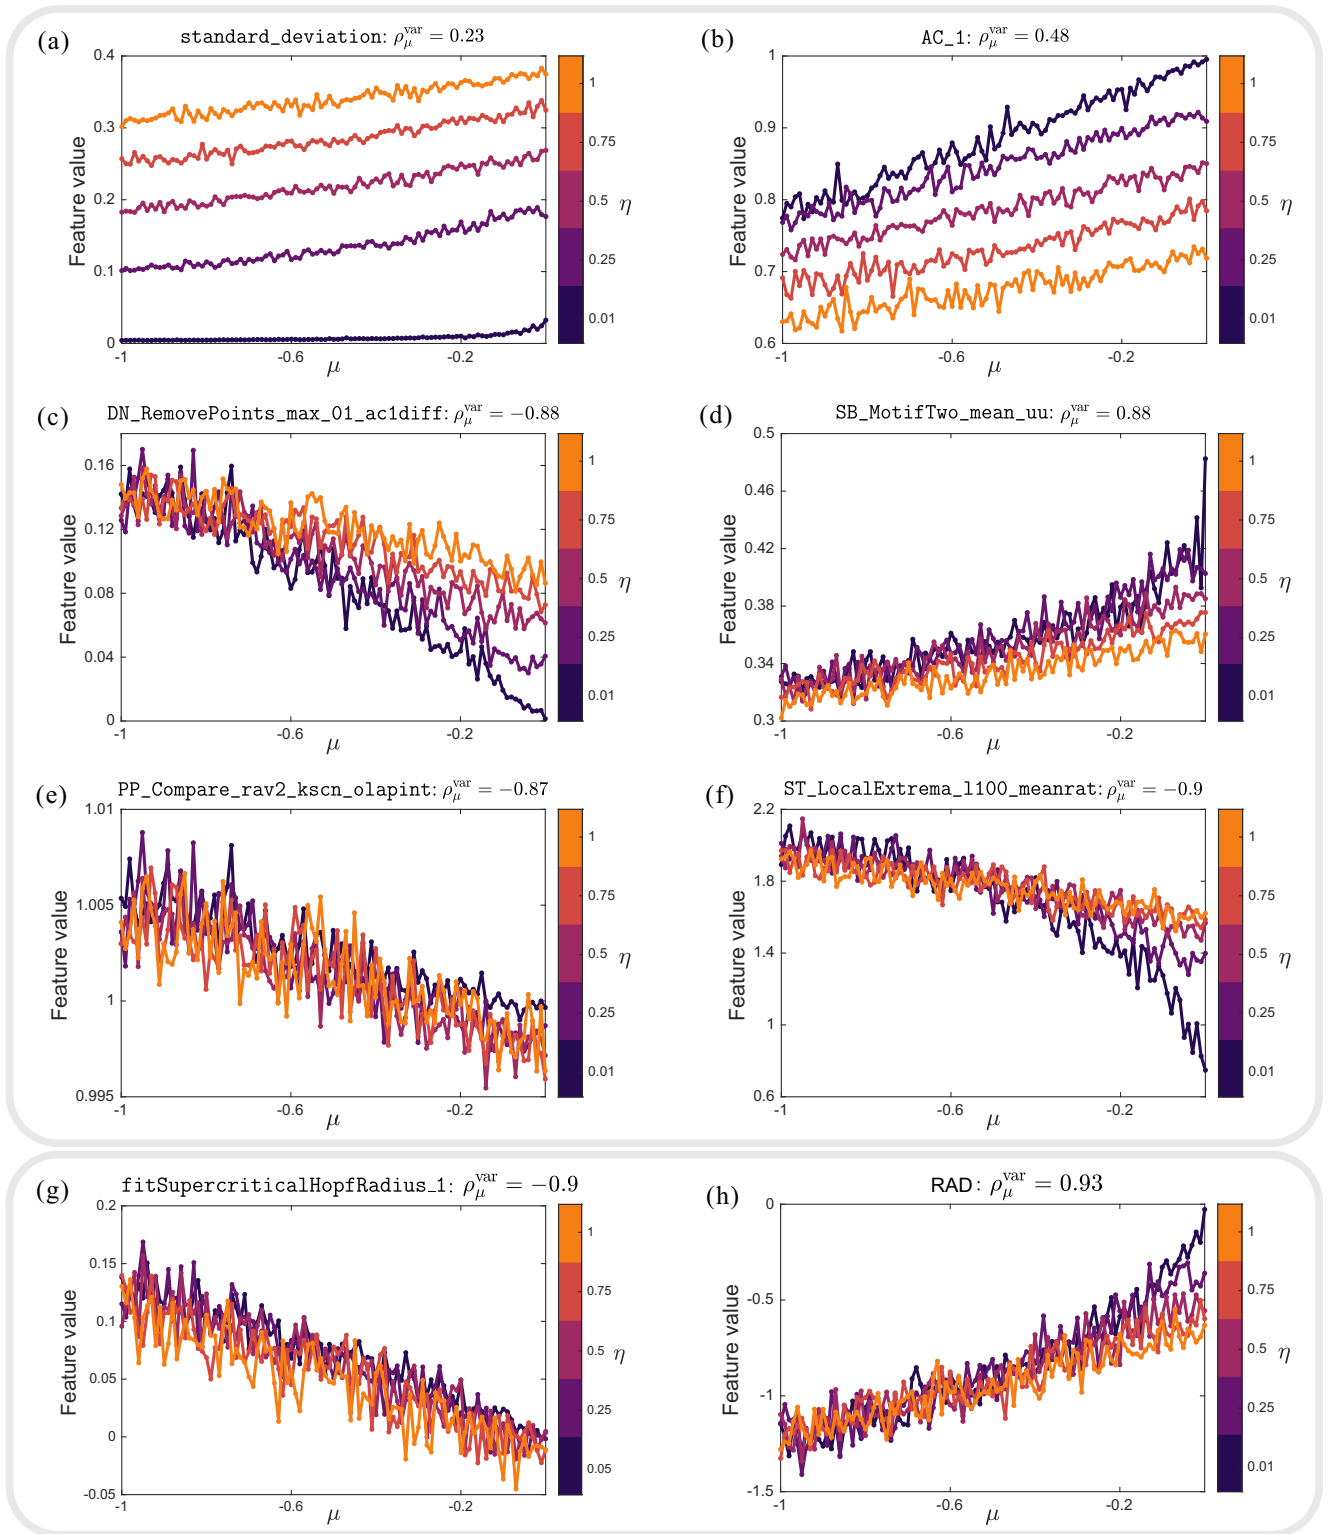


FIG. 4. Conventional features for tracking the DTC perform well only in the fixed-noise setting, but we uncover new features that robustly track the DTC across different noise amplitudes. Feature values are plotted against μ for five values of the noise amplitude $\eta = 0.01, 0.25, 0.5, 0.75, 1$. Two of the top fixed-noise features, (a) standard deviation and (b) lag-1 autocorrelation, are highly correlated to μ for any given value of η (the fixed-noise case), but are poorly correlated across multiple η values (the variable-noise case). (c)–(f) The highest-performing time-series features in the variable-noise case are highly correlated to μ both within and across η values, demonstrating robust tracking of the DTC despite unknown noise amplitude. (g) A new feature introduced here, `fitSupercriticalHopfRadius_1`, distills the key algorithmic elements of the top-performing features. However, the curve-fitting algorithm is unstable for $\eta < 0.05$, so only $\eta = 0.05$ is shown for this feature. (h) A second new feature, `RAD`, uses elementary time-series operations to yield a computationally efficient and numerically stable estimate of the DTC. Each feature is summarized in Table I and detailed in Sec. III C.

TABLE I. Summary of top-performing time-series features from the *hctsa* feature library. The magnitude of ρ_μ^{fix} and ρ_μ^{var} indicate how well each feature can track the distance to criticality in the fixed-noise and variable-noise settings, respectively. Our new features are `fitSupercriticalHopfRadius_1` and the RAD, `CR_RAD_1`.

Feature name	ρ_μ^{fix}	ρ_μ^{var}	Description
<code>standard_deviation</code>	0.98	0.23	Sample standard deviation
<code>AC_1</code>	0.97	0.48	Lag-1 autocorrelation
<code>DN_RemovePoints_max_01_ac1diff</code>	0.94	-0.88	Change in lag-1 autocorrelation from removing the largest 10% of values
<code>SB_MotifTwo_mean_uu</code>	0.93	0.88	Proportion of pairs of consecutive values that are both above the mean
<code>PP_Compare_rav2_kscn_olapint</code>	0.88	-0.87	Change in distribution after two-sample moving-average smoothing
<code>ST_LocalExtrema_1100_meanrat</code>	0.91	-0.90	Ratio between the average maximum and minimum in 100-sample windows
<code>fitSupercriticalHopfRadius_1</code>	0.92	-0.90	Potential function curve fit to a kernel-density estimate of the distribution
<code>RAD</code>	0.93	0.93	Product of the spread of differences and tailedness of the distribution

toward the origin. This change in the shape of the potential results in dynamics with slower fluctuations (increased autocorrelation) and diffusion over a larger domain [increased standard deviation; cf. orange distributions in Figs. 5(a)–5(d)]. However, Figs. 5(a)–5(d) also show that η can make the relationship between the potential and autocorrelation or standard deviation ambiguous: The time-series autocorrelation takes a similar value for $\mu = -2, \eta = 0.5$ [Fig. 5(a)] as for $\mu = -0.1, \eta = 1.5$ [Fig. 5(d)], and the distribution has a similar standard deviation for

$\mu = -2, \eta = 1.5$ [Fig. 5(b)] and $\mu = -0.1, \eta = 0.5$ [Fig. 5(c)]. While both the potential function and driving noise determine how a system evolves over time, we can see that a high-performing algorithm for estimating μ in the presence of variable noise should target a property that is specifically affected by μ : the shape of the potential function $V(x; \mu)$.

2. The invariant density $p(x)$

Under the conceptual framing provided by the potential formulation, we next aimed to investigate whether the

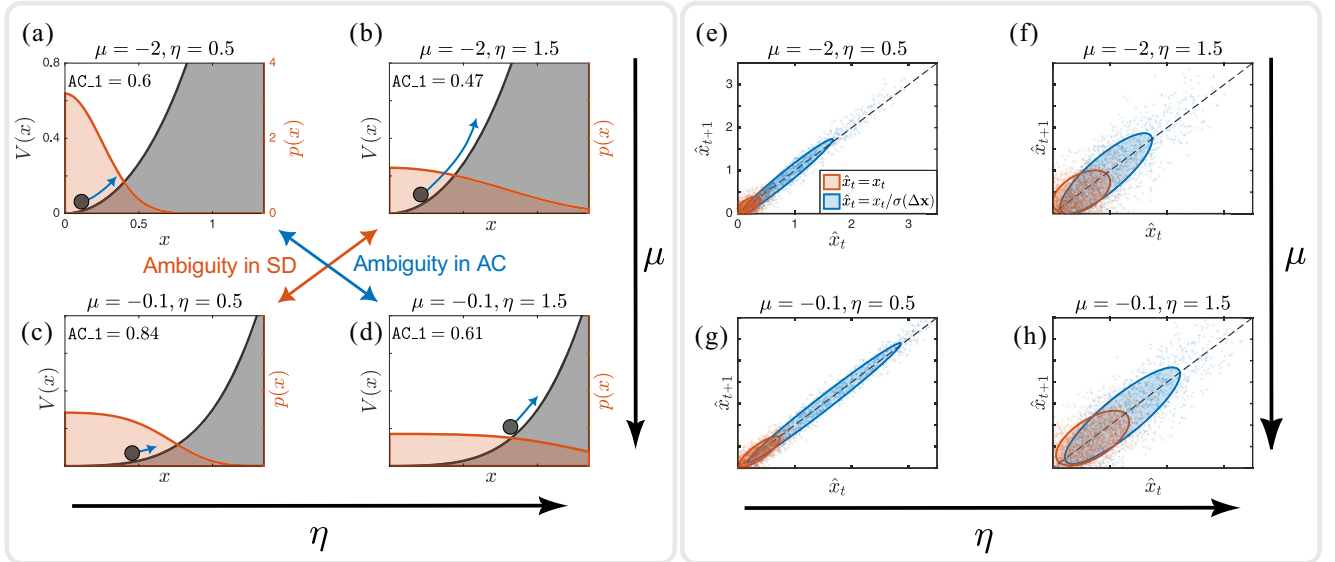


FIG. 5. Rescaling time-series values by the spread of differenced values corrects for the confounding effect of a variable-noise amplitude. (a)–(d) We plot the potential function $V(x)$ (determined by the control parameter μ shown black) and distribution of values $p(x)$ (orange) for four combinations of μ and η : (a) $\mu = -2, \eta = 0.5$, (b) $\mu = -2, \eta = 1.5$, (c) $\mu = -0.1, \eta = 0.5$, and (d) $\mu = -0.1, \eta = 1.5$. The control parameter μ defines the DTC and determines the potential function $V(x)$. The flatness of $V(x)$ gives rise to critical slowing-down that, in combination with the noise amplitude η , determines $p(x)$ and the lag-1 autocorrelation (annotated as AC_1, top right). Both $p(x)$ and lag-1 autocorrelation values give ambiguous estimates of the DTC: Different pairs of μ and η can result in the same feature value. (e)–(h) Scatters of (x_t, x_{t+1}) (orange) and the rescaled $[x_t/\sigma(\Delta \mathbf{x}), x_{t+1}/\sigma(\Delta \mathbf{x})]$ (blue) for the same four parameter settings as in (a)–(d): (e) $\mu = -2, \eta = 0.5$, (f) $\mu = -2, \eta = 1.5$, (g) $\mu = -0.1, \eta = 0.5$, and (h) $\mu = -0.1, \eta = 1.5$. After rescaling by the standard deviation in the $x_{t+1} - x_t$ direction $\sigma(\Delta \mathbf{x})$ (described in Sec. III D), the distribution of $x_t + x_{t+1}$ becomes less sensitive to changes in the noise amplitude. Annotated ellipses indicate the covariance of (x_t, x_{t+1}) (orange), which is strongly affected by changes in η , but becomes more characteristic of μ after rescaling by $\sigma(\Delta \mathbf{x})$ (blue).

top-performing features were able to extract a noise-robust estimate of μ from time-series data by estimating the shape of the potential $V(x; \mu)$. As described above, all top-performing features measure properties of the distribution, suggesting that this step may be relevant to robustly estimating the DTC (see Sec. III B). Following this connection, we investigated the invariant density $p(x)$: the probability density to which the distribution of values from a stationary system will converge over time. The invariant density is valuable for describing bifurcations when the abruptness of the qualitative change that occurs in the deterministic system is destroyed or smoothed by noise [58]. For a stationary system, the invariant density can be derived from the potential function and the noise term by taking the Fokker-Planck equation in the limit of infinite time [81]. For the system studied here, Eq. (1), the invariant density has the form

$$p(x; \mu, \eta) = A \exp\left(\frac{-2V(x; \mu)}{\eta^2}\right), \quad x \geq 0, \quad (3)$$

which is normalized to unit probability mass by A [58]. For fixed η , the potential can be inferred from the invariant density as

$$V(x; \mu) = -\frac{1}{2}\eta^2[\ln p(x) - \ln A], \quad (4)$$

where it is straightforward to obtain an estimate of the probability density $p(x)$ from measured data (e.g., using a kernel estimator [87]). Since Eq. (2), and therefore Eq. (4), depend on μ , time-series features that measure properties of the distribution of values in the time series (such as mean, standard deviation, and skewness) are able to provide accurate estimates of the DTC in the fixed- η case (as we verified in Sec. III A). However, as Eq. (4) depends strongly on η , features measuring the distribution of time-series values are highly sensitive to noise and thus have low ρ_μ^{var} (as we found in Sec. III B).

We now seek to incorporate the second algorithmic component of the top-performing features, the spread of differences, into our expression for the potential function, in the hope that it may eliminate the noise amplitude η from the invariant density and yield a noise-insensitive estimate of μ .

3. The spread of differences, $\sigma(\Delta \mathbf{x})$

In addition to the distribution of time-series values, the time-series features with the highest ρ_μ^{var} also indirectly measure the spread of the incrementally differenced time series $\Delta \mathbf{x}$, where Δ is the first difference operator and $\Delta x_t = x_t - x_{t-1}$ [88]. We write this spread here as $\sigma(\Delta \mathbf{x})$, where σ is the standard deviation (see Sec. A for further details). Considering the drift and diffusion terms of Eq. (1), we see that the standard deviation of increments of the Wiener process scale with $\sqrt{\Delta t}$ [89], whereas

increments of $(\mu x - x^3)dt$ scale with Δt . Hence, for a small time step Δt the standard deviation of the differenced time series is dominated by the stochastic component ηdW . In the short-time-step limit $\Delta t \rightarrow 0$, such that $\sqrt{\Delta t} \gg \Delta t$, we obtain an approximation for inferring η :

$$\sigma(\Delta \mathbf{x}) \approx \eta \sqrt{\Delta t}, \quad (5)$$

where $\sigma(\Delta \mathbf{x})$ is the standard deviation of the incrementally differenced time series $\Delta \mathbf{x}$.

The top-performing variable-noise features all involve measuring this dynamical quantity $\sigma(\Delta \mathbf{x})$ that is informative of η . This insight allows us to understand how they may be able to eliminate the noise amplitude term from the invariant density to more accurately estimate the DTC. Incorporating this empirical estimate of η from Eq. (5) into our expression for the potential in terms of the invariant density, Eq. (4), yields

$$V(x; \mu) \approx -\frac{\sigma^2(\Delta \mathbf{x})}{2\Delta t} [\ln p(x) - \ln A], \quad (6)$$

where we assumed a fixed sampling period Δt .

To test Eq. (6), we used it to infer the DTC in the fixed-noise and variable-noise cases studied here. For the system given by Eq. (1), we can infer $V(x)\Delta t$, and therefore the DTC, using two properties that are easily measured from time-series data: the probability density $p(x)$ and the spread of differences $\sigma(\Delta \mathbf{x})$. The algorithmic procedure for inferring the DTC from data using Eq. (6) is as follows: (i) produce a kernel-density estimate for $p(x)$; (ii) calculate $\sigma(\Delta \mathbf{x})$ using the standard deviation and the first difference operator [88]; (iii) transform the density to a potential using Eq. (6); (iv) curve fit Eq. (2) to the estimated potential; (v) infer μ by calculating the value of the second derivative at $x=0$. We give this procedure the feature name `fitSupercriticalHopfRadius_1` and show how it varies with both μ and η in Fig. 4(g).

Since `fitSupercriticalHopfRadius_1` varies strongly with the control parameter μ but minimally with the noise amplitude η , we find that it tracks the DTC in the variable-noise case far more strongly than conventional fixed-noise metrics, yielding scores that are on par with the top-performing *hctsa* features (see Sec. III B): $\rho_\mu^{\text{fix}} = 0.92$ and $\rho_\mu^{\text{var}} = 0.90$ (see Supplemental Material Table S1 [84] for a complete list of scores for both new features and *hctsa* features). The strong performance of `fitSupercriticalHopfRadius_1` suggests that Eq. (6) is indeed capturing key algorithmic principles relevant to robustly tracking the DTC. However, the top *hctsa* features (detailed in Appendix A) are able to infer the DTC with disparate and often simple algorithms that do not rely on the computationally expensive steps of density estimation and curve fitting, and do not require explicitly knowing the form of the potential or the noise process. Furthermore, we find that

`fitSupercriticalHopfRadius_1` suffers from numerical instability and produces noisy values when the noise amplitude is low since the distribution is tightly concentrated around $x \approx 0$; in our simulations, this instability occurred only at $\eta = 0.01$. Armed with a theoretical explanation of how time-series features are able to estimate the DTC in the presence of noise, we next aim to develop a simple, efficient, stable, and generic algorithm that implements the potential-based inference of the DTC described above.

D. Rescaled autodensity: A new noise-robust metric of the distance to criticality

In this section, we develop and test a new feature that directly implements the key algorithmic steps underpinning the success of the top-performing *hctsa* features, allowing us to efficiently estimate the DTC from data. We first detail a simplified approach to inferring the potential function $V(x)$ [cf. Eq. (6)] that eliminates η by rescaling time-series values by the standard deviation of incremental time-series differences, $\sigma(\Delta\mathbf{x})$. We then describe how the density of points in an (x_t, x_{t+1}) time-delay embedding can be used to infer the DTC from the noise-robust rescaled potential. Finally, we introduce the RAD, a new time-series feature that uses elementary time-series operations to implement these key algorithmic principles.

1. Rescaled potential

We first notice that when the system is close to the origin $x \ll 1$, the quadratic term in Eq. (2) $-\mu x^2/2$ dominates the quartic term $x^4/4$. Since our system is driven only by noise and has no extreme jumps or perturbations, it is naturally confined to a region close to the origin by the “steep walls” of the quartic term [see Figs. 5(a)–5(d)]. This behavior can be observed from the Taylor expansion of the invariant density Eq. (3) given by

$$\tilde{p}(x; \mu, \eta) = B \left(1 + \mu \frac{x^2}{\eta^2} + \frac{\mu^2 x^4}{2 \eta^4} - \frac{x^4}{2\eta^2} + \mathcal{O}(x^6) \right), \quad (7)$$

where B is a constant that normalizes $\tilde{p}(x)$ to unit probability mass. Note, however, that this behavior breaks down when η becomes too large (the noise drives the system to extreme values where the quartic term dominates) or μ becomes too small (the quartic term dominates as the quadratic term is silenced). Under the approximation that the system tends to reside close to equilibrium $x \ll 1$, we can write a simplified, rescaled potential as

$$\hat{V}(\hat{x}; \mu) = -\frac{\mu \hat{x}^2}{2} = -\frac{\ln[p(\hat{x})] - c}{2\Delta t}, \quad (8)$$

where $\hat{x} \geq 0$, c is a constant, and $\hat{x} = x/[\sigma(\Delta\mathbf{x})] \approx x/(\eta\Delta t)$. This approximate potential \hat{V} provides an algorithmically

simpler way to estimate the DTC: We first rescale the system with $\hat{x} = x/[\sigma(\Delta\mathbf{x})]$ (aiming to eliminate the contribution of η) and then infer the potential of the rescaled system using the distribution of the rescaled values (to estimate μ). Unlike the theory described above, in Sec. III C, this simplified, highly approximate approach is not specific to the potential function of our model system.

2. The autodensity

Having identified a simpler method for “sensing” the underlying potential function governing the deterministic dynamics, by rescaling the system with the standard deviation of incremental time-series differences, we now develop a simple algorithm for estimating the DTC from the inferred potential. We notice that the key attributes for robustly inferring the DTC—the invariant density and the spread of differences—can both be measured from the distribution of time-series values in an (x_t, x_{t+1}) time-delay embedding [90], as illustrated in Figs. 5(e)–5(h). For the purposes of naming our feature, we refer to this two-dimensional distribution in (x_t, x_{t+1}) as the lag-1 autodensity (in analogy to autocorrelation). This autodensity captures the two properties we have found to be key for robustly inferring the DTC: the spread of differences $\sigma(\Delta\mathbf{x})$ and the distribution $p(x)$. First, in a two-dimensional embedding space (x_t, x_{t+1}) , the factor $\sigma(\Delta\mathbf{x})$ is proportional to the standard deviation of shortest distances from $x_t = x_{t+1}$ or $x_{t+1} - x_t$. Second, when the sampling period is small $\Delta t \ll 1$, the distribution $p(x)$ is well approximated by the linear projection of x_t values onto the line defined by $x_t = x_{t+1}$ proportional to the distribution of the quantity $x_t + x_{t+1}$. The approximation given by Eq. (8) can then be implemented by rescaling x_t and x_{t+1} with the width of the autodensity in the $x_{t+1} - x_t$ direction. As illustrated in Figs. 5(e) and 5(f), rescaling the time-series values produces a distribution that is minimally sensitive to changes in the noise amplitude.

3. The rescaled autodensity

We now encapsulate the key algorithmic steps described above into a new time-series feature, RAD, that aims to robustly infer the DTC of noisy systems. It begins by partitioning the time-series values about a threshold that is insensitive to changes in extreme densities. The median (\tilde{x}) is a suitable choice for this threshold, partitioning the time series into two sets: time-series values above the median

$$U = \{x_t : x_t \geq \tilde{x}\}, \quad (9)$$

and the time-series values below the median

$$L = \{x_t : x_t < \tilde{x}\}. \quad (10)$$

RAD then aims to summarize of the shape of the invariant density while avoiding the need for the curve fitting

(as required by `fitSupercriticalHopfRadius_1`; cf. Sec. III C). For a simple measurement of the shape of the probability density, we quantify the tailedness as the difference in the average density between the upper and lower partitions. Here we define the average density in each partition as the ratio of the probability mass contained in the partition and the width of the partition measured by the standard deviation. Noting that the two partitions are split by the median, and therefore have an equal probability mass of 0.5, we can capture the average density in a partition as the inverse of the standard deviation $1/\sigma$. Rescaling the difference between the average densities of the upper and lower partitions (which measures the shape of the distribution, in particular, the higher-order moments) by $\sigma(\Delta\mathbf{x})$ completes our RAD feature, which is then given by

$$f_{\text{RAD}} = \sigma(\Delta\mathbf{x}) \left[\frac{1}{\sigma(U)} - \frac{1}{\sigma(L)} \right]. \quad (11)$$

RAD takes negative values far from criticality and approaches zero when a system is close to the critical point. We also define the centered RAD denoted as cRAD, which is applicable to nonradial data:

$$f_{\text{cRAD}} = f_{\text{RAD}}(|x - \tilde{x}|), \quad (12)$$

where \tilde{x} is the median value of \mathbf{x} and $|\cdot|$ is the absolute value. To verify the DTC-tracking performance of RAD, we applied it to the fixed-noise and variable-noise cases analyzed above, as shown in Fig. 4(h). In the variable-noise case, RAD outperformed all other *hctsa* features, with $\rho_{\mu}^{\text{fix}} = 0.93$ (see the RAD feature in Supplemental Material Table S1 [84]). RAD also exhibited competitive performance relative to the top features in the fixed-noise setting, with $\rho_{\mu}^{\text{var}} = 0.93$. Furthermore, we verified that RAD outperforms conventional metrics in a range of other systems with various normal forms and noise processes (see Sec. S2 in Supplemental Material [84]). RAD is a straightforward and transparent algorithm that behaves as a reliable indicator of DTC in noisy systems. MATLAB, JULIA, and PYTHON implementations of RAD are provided in an accompanying code repository [91] and are also included in the *hctsa* time-series feature library [68].

E. Structure-function organization across the mouse visual cortex

We next aimed to test the performance of RAD in a real-world setting using electrophysiological data, alongside conventional metrics of tracking the DTC. Real-world complex systems, such as the brain, are thought to exploit near-critical states for computational advantage [7,9,92,93], with near-critical systems exhibiting enhanced dynamic range [92,94,95], input separation and sensitivity, information-storage capacity, and information-transfer capabilities [5,41,96]. In neural systems, which benefit

from such advantages, evidence for criticality has been found at multiple scales, from single neurons [33,97,98] to neuronal spike trains [99–106], local field potentials [107–109], as well as magnetoencephalography [110] and functional magnetic resonance imaging time series [111].

Numerous studies have suggested that, despite the commonly reported benefits of criticality, the brain tends to operate in a subcritical state [112–114], often associated with focused attention [115,116]. However, other research suggests that the brain fluctuates around a critical threshold [9,106]. A more comprehensive perspective, considering the diversity of brain regions, suggests the advantages of integrating regions with varying dynamical states, combining those closer to criticality with others that are more subcritical [76,95,117]. Many theories of neural criticality then predict that regions further along the cortical hierarchy (i.e., regions responsible for higher-order cognitive functions such as the integration of complex stimulus features, conscious perception, or decision-making [118]) require longer timescales for integrating information [71,72,74,119], and are thus expected to sit closer to criticality [7,76,117]. This hypothesis has been supported, in part and indirectly, by a corresponding hierarchical variation in the excitation-to-inhibition ratio [76,120] and other structural features [73,121], structural and functional connectivity [122,123], as well as intrinsic timescales [72,74] across cortical areas. While much work has analyzed the theoretical and anatomical basis for tuning the DTC of neural systems [7,124], the presence of variable stochastic drives across brain regions—perhaps arising from differences in connectivity to external regions [125], nearby vasculature [126], or heterogeneous cytoarchitectures and their influence on high-frequency dynamics [127]—confound the inference of the DTC in real-world macroscale recordings of neural activity using conventional measures like autocorrelation and variance.

Based on the hypothesis that the proximity of a cortical area to criticality increases along the visual hierarchy [7,76,117], and that brain regions may be subject to different levels of dynamical noise, here we aimed to compare how well RAD tracks the hierarchical level—measured independently using anatomical data [128,129]—against conventional measures: lag-1 autocorrelation and standard deviation. We used experimental electrophysiological recordings of the mouse brain from the Allen Neuropixels Visual Behavior dataset [130], which provides local field potential (LFP) data from approximately 20 channels (between 15 and 24 for any given mouse) in each of six regions of the mouse visual cortex [shown schematically in Fig. 6(a)]. We selected recording sessions based on quality metrics, extracted LFP data from periods where mice were viewing a static gray screen, applied a (1–20) Hz bandpass filter, and finally decimated the signals to a sampling rate of 125 Hz, yielding approximately 20 time series (approximately 5 min, or approximately 37 500 samples long) for each of six visual cortical regions over 39 sessions

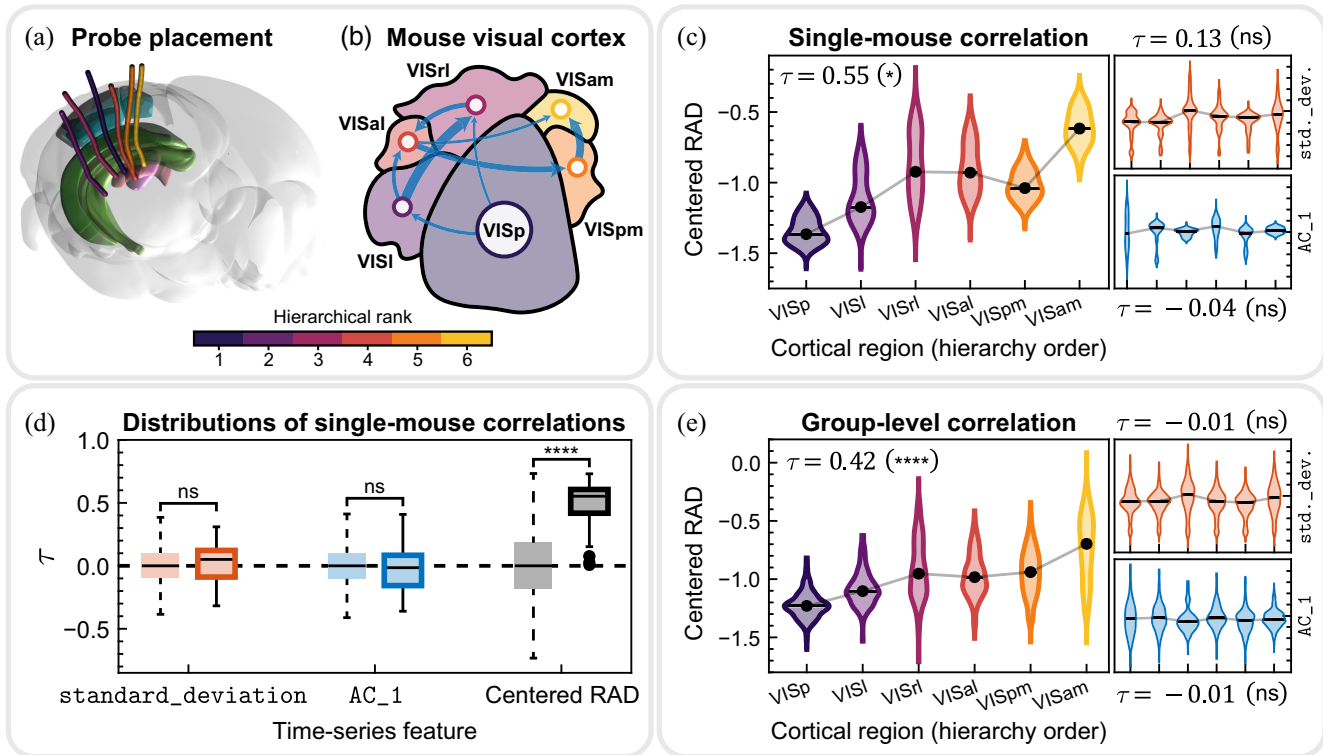


FIG. 6. RAD tracks the anatomical hierarchy of mouse visual cortical regions from neural electrophysiology data. (a) Six Neuropixels probes (lines, black to yellow) in the mouse brain and the key areas they intersect (visual cortex, blue; hippocampus, green; thalamus, pink). (b) Six visual cortical regions and their hierarchy ranks (black to yellow) shown with anatomical similarities from Siegle *et al.* [128] (blue arrows, width weighted by similarity, directed from lower-order to higher-order areas). (c) The distribution of centered RAD over $n \in [15, 24]$ channels for each of the six cortical regions in a single representative mouse (session 1065908084). The strong Kendall's τ correlation of centered RAD to the hierarchy score ($\tau = 0.55$, $p = 0.03$; *) is annotated, and is equal to the group median; see (d). Connected dots show the median centered RAD, and regions are ordered based on the anatomical hierarchy shown in (b). Outside right are similar plots for standard deviation (red, top right; $\tau = 0.13$, $p = 0.4$) and lag-1 autocorrelation AC_1 (blue, bottom right; $\tau = -0.04$, $p = 0.7$), both showing weak and nonsignificant (ns) correlations. (d) The distribution over mice ($n = 39$) of Kendall's τ correlations, as depicted by box plots in (b), for standard deviation (red), lag-1 autocorrelation (blue), and centered RAD (gray). Bars show median values, boxes show the interquartile range, dots show outliers ($>1.5 \times IQR$ from the nearest quartile), and whiskers show the range of nonoutlier values. p -values for Mann-Whitney U tests between the correlations for each feature and their corresponding null distributions [shown by dashed box plots, over $n = 39 \times 10^6$ values of τ ; (see main text)] are annotated. Single-mouse correlations of centered RAD to the hierarchical level are highly significant, with $p < 10^{-18}$ (**** signifies $p < 10^{-4}$), but are not significant for both standard deviation ($p = 0.3$) and lag-1 autocorrelation ($p = 0.5$). (e) The distribution of centered RAD pooled across all mice and channels (yielding $n \approx 750$ data points for each region). The group-level Kendall's τ computed on this pooled data is annotated alongside an estimated p -value (see main text). Even at a group level, the correlation of centered RAD to the hierarchy rank is strong and highly significant ($\tau = 0.42$, $p < 10^{-6}$), unlike both standard deviation (red, top right; $\tau = -0.01$, $p = 0.8$) and lag-1 autocorrelation (blue, bottom right; $\tau = -0.01$, $p = 0.5$).

(39 mice; see Appendix B). We then calculated RAD, lag-1 autocorrelation, and standard deviation for each time series. Note that we used centered RAD, Eq. (12), appropriate for these data with negative values. To rank regions according to their position in the anatomical hierarchy, we used data on the hierarchical level from Siegle *et al.* [128], who recomputed the connectome-based scores from Harris *et al.* [129] for the visual regions of the Allen Neuropixels dataset [see Fig. 6(b)]. We performed a one-way analysis of variance on feature values (averaged across channels for a given probe) against cortical regions (using data pooled across all mice) to determine if there were significant differences

in the distribution of feature values between regions. We identified significant effects for each of centered RAD [$F(5, 190) = 66$, $p < 10^{-38}$], standard deviation [$F(5, 190) = 14$, $p < 10^{-11}$], and lag-1 autocorrelation [$F(5, 190) = 4$, $p = 1.5 \times 10^{-3}$], encouraging more detailed analysis of the variation of these features across the hierarchy.

We then sought to determine, at the level of individual mice, the strength of the relationship between the values of a given feature (computed across LFP channels) and the hierarchical ordering of the cortical areas of each channel. We first calculated Kendall's τ (a tie-robust correlation

coefficient) between feature values and hierarchical ranks across all channels, as shown for a single representative mouse in Fig. 6(c). We used a permutation-based procedure to estimate a p -value for each computed τ statistic, generating an empirical null distribution for each feature by shuffling, for each mouse, the hierarchical ranks assigned to each region (ensuring that channels from the same region retained the same rank, thereby accounting for spatial correlation between nearby channels). After calculating a τ coefficient for each of 10^6 independently shuffled datasets, p -values were estimated as the proportion of surrogate correlations, τ that were larger in magnitude than the τ statistic of the real data. We found that regions lower in the visual hierarchy (e.g., the primary visual area, “VISp”) have a lower value of centered RAD—corresponding to a greater DTC—than higher regions (e.g., the anteromedial visual area, “VISam”). To assess this relationship across mice, we performed a Mann-Whitney U test between the distributions of correlations for the original and shuffled datasets. As shown in Fig. 6(d), neither standard deviation nor lag-1 autocorrelation had single-mouse correlation coefficients that were significantly different from their corresponding null distributions (at a threshold of $p = 0.05$). By contrast, centered RAD achieved strikingly strong correlations for most mice (a median $\tau = 0.55$), and was highly significant ($p < 10^{-18}$). That this signature was detectable by the noise-robust RAD but not lag-1 autocorrelation or standard deviation was consistent with the presence of variable-noise levels across regions, masking the effects of criticality from traditional features.

Our results above provide evidence consistent with the hypothesis that, in a given mouse, brain regions higher in the visual hierarchy are closer to criticality. We next aimed to test whether the result also held across mice, i.e., that variability due to the hierarchical variation detectable by RAD was stronger than intermouse variability. To achieve this, we pooled the data from all channels in all mice to calculate a group-level τ coefficient for each feature [shown in Fig. 6(e)]. We estimated the statistical significance of group-level correlations using a similar permutation-based procedure as above (i.e., calculating the proportion of 10^6 shuffles that had a correlation greater than the true measured value for a feature). Remarkably, centered RAD remained highly correlated to the visual hierarchy even at the group level ($\tau = 0.42$, $p < 10^{-6}$), whereas standard deviation ($\tau = -0.01$, $p = 0.8$) and lag-1 autocorrelation ($\tau = -0.01$, $p = 0.5$) did not detect any significant relationship, as per the individual-level analysis. Our results suggested that higher-order cortical regions exhibited dynamics consistent with being closer to criticality [7, 117], even though they may have been influenced by varied levels of noise due to external input, cellular composition, laminarity, or the presence of vasculature. Together with the hypothesis that sensory regions are differentiated by their proximity to criticality, these results

provide strong evidence that RAD can be a more reliable metric of the DTC than conventional metrics in noisy, real-world systems.

IV. DISCUSSION

This work addressed the challenge of estimating the distance to a critical point DTC in the presence of an uncertain, and potentially variable, confounding noise amplitude. While many studies have tackled the challenge of describing how noise disrupts conventional metrics of the DTC [13, 14, 47, 58] or intensifies low-amplitude perturbations by stochastic resonance [131], to our knowledge no work has identified a new feature that is insensitive to possible changes in noise levels and demonstrated its efficacy in a real-world system. Moreover, our work highlighted the ability of our novel data-driven methodology to motivate new theory and algorithmic implementations relevant to working with real (finite and noisy) time series. Specifically, by comparing the performance of thousands of time-series analysis features, we identified analysis methods that were able to robustly track the DTC in systems with variable levels of dynamical noise, but which have not been applied to this problem in the past. By analyzing the algorithmic steps underlying top-performing time-series features, we developed a deeper theoretical understanding that enabled us to reverse engineer a novel, efficient, and noise-robust index of the DTC that outperformed all other *hctsa* features on this problem: the RAD.

Our data-driven approach to solving theory-based problems is highly flexible and could be extended to many real-world problems, particularly those involving short, noisy time-series data. The approach involves first simulating a known dynamical mechanism (to generate time-series data with a known structure) and then searching across a sufficiently large and comprehensive set of candidate time-series features for those that can best recover the underlying structure. This highly comparative methodology has been used previously by Fulcher *et al.* [67] to find statistical estimators of the scaling exponent of self-affine time series, and the Lyapunov exponent of Logistic Map time series; problems for which high-performing features were known to exist within the candidate feature set. This is also the case for the first fixed-noise setting investigated here, where our data-driven approach recapitulated the strong performance of conventional criticality metrics related to autocorrelation and the distribution of values. However, here we also extended it to a new setting, the variable-noise setting, in which it was not known whether any features would exhibit strong performance. We showed that conventional metrics with high ρ_μ^{fix} (which strongly correlate with the DTC in the fixed-noise setting) perform poorly in this variable-noise setting, with a low ρ_μ^{var} (Sec. III B). But, surprisingly, we identified several high- ρ_μ^{var} features that can track DTC in the presence of confounding variations in noise amplitude.

Having identified noise-robust features of DTC from *hctsa*, we then showed how our data-driven methodology can motivate the development of new theory and understanding. While a list of top-performing features is already a valuable resource for addressing the problem of inferring the DTC—and brings about the possibility of using an ensemble of metrics for a more accurate estimate—we also sought practical insight regarding these top-performing metrics, the problem itself, and the broader context of criticality in noisy systems. In this capacity, our approach sits alongside recent efforts to generate practical understanding from data using interpretable machine-learning algorithms [132,133] or to solve analytical problems using artificial intelligence [134]. Unlike existing approaches, our methodology uses a simple search across a comprehensive library of transparent algorithms drawn from existing literature, automatically flagging those that are relevant to a given problem. Studying the highlighted algorithms can then yield new theoretical insight into the problem at hand. In this work, noticing key algorithmic similarities in the top-performing features motivated us to develop a theoretical account of how these algorithmic steps were able to track DTC so successfully. The resulting theory—invoking a potential formulation Eq. (2), and the corresponding expression for the invariant density Eq. (3)—was used to formulate a new high-performing time-series feature, RAD, that robustly infers the shape of the potential function, which depends only on the DTC, by measuring the time-series distribution after rescaling values with the spread of differences (see Sec. III D). RAD performs well in both noise settings, having a higher correlation to the DTC than all other features in the variable-noise setting, as well as competitive performance (relative to conventional metrics) in the fixed-noise setting. As such, RAD requires less prior knowledge of a system to produce an accurate estimate of the DTC, making it a practical statistic for tracking the DTC in real-world settings. Given the increasingly broad and detailed datasets being generated across scientific domains, this work thus demonstrates an ability to derive theoretical insight and develop practical analytic tools via the broad algorithmic comparison enabled by large algorithmic libraries like *hctsa* [68]. Our approach serves as a model for using wide methodological comparison to tackle similar problems that aim to develop new theory for bridging dynamical mechanisms with the statistical properties, such as criticality, that are most sensitive to the theoretical structures of interest.

No previous studies have tackled the challenge of developing noise-robust indicators of the DTC by distilling theory from data-driven exploration. However, the algorithmic components of RAD, including the spread of differences and a measurement of asymmetry in the distribution Eq. (11), share similarities with existing metrics related to criticality. For instance, the spread of differences (see Sec. III C) has been used to anticipate

critical transitions in cryptocurrency markets [135], whereas skewness and related properties of the distribution (see Sec. III C) have been used to mark abrupt changes in ecosystems [136] and climates [137]. Crucially, neither the spread of differences nor the distribution alone give noise-robust features. As we highlighted here, the confounding influence of the noise amplitude can be eliminated only by carefully combining these two properties (see Sec. III D). Furthermore, estimating the DTC typically requires features to be precisely calibrated against the control parameter and the noise amplitude with repeated observation on a system-by-system basis. Methods for calibrating indicators of criticality, which have been addressed elsewhere and are not the focus of this work, include learning simple thresholds [18], performing linear regression [138], or training supervised machine-learning models [139] to predict critical phenomena from labeled data—RAD can be readily dropped into these existing pipelines, alongside the conventional metrics of criticality. We briefly illustrated, using linear regression and our simulated dataset, how the top variable-noise features can be calibrated to make concrete predictions of the DTC in Sec. S1(b) of Supplemental Material [84]. RAD still requires calibration to the control parameter, as for typical indicators of the DTC, but there is less need to recalibrate for changes or uncertainty in noise. Moreover, RAD is superior even in contexts that do not require concrete estimates of the DTC; we showed in Sec. III E that being robust to noise allows RAD to more accurately recover the relative positioning of brain regions in the mouse visual hierarchy. RAD also has many other advantages over other existing critical indicators that have been applied in noisy settings, including (i) it does not require perturbing the system [140], (ii) it operates on univariate time series [51], and (iii) it performs well on short time series (but assumes a small sampling period). Many real-world systems are corrupted by noise with an unknown strength, can be recorded only for short periods, and cannot be measured in their full multivariate complexity. As such, RAD may improve accuracy in the wide range of tasks that involve quantitatively anticipating criticality from time-series data.

We verified the real-world utility of RAD by applying it to a dataset of electrophysiological recordings from the mouse visual cortex, where it outperformed conventional metrics of the DTC in tracking the anatomical hierarchy of visual cortical regions (in Sec. III E). A long-standing hypothesis on cortical organization is that higher-order regions are closer to criticality, facilitating the integration of multiple external signals or stimulus features [7,76,117], but this hypothesis has lacked concrete experimental evidence. Our finding that RAD increases along the mouse visual hierarchy offers a mechanistic explanation for the escalating gradient of slow fluctuations along the visual hierarchy as being caused by critical slowing-down, with regions higher in the hierarchy exhibiting a smaller DTC.

Furthermore, our results suggest that in real-world neuronal systems, noise is better represented heterogeneously, with variable rather than fixed intensity. Variable noise renders conventional time-series metrics of criticality, which are highly sensitive to variations in noise level between brain regions, blind to variation in the DTC. Future work could aim to extend our analysis of heterogeneity in the DTC to the broader cortex, incorporating LFP recordings from other cortical areas or whole-cortex data using imaging modalities with a wider field of view [141]. Moreover, changes and disruptions to the DTC have been implicated in broader neural contexts [142]: for insomnia disorder and the transition to sleep [17], autism spectrum disorder [143], and epilepsy at the transition to seizure [18,19]. In these instances, employing RAD could also present a pertinent and innovative approach. For seizure anticipation especially, simple but reliable indicators that can be applied to noisy, noninvasive recordings are essential for building practical monitoring devices [18,19]. Having been formulated for robustness against dynamical noise (in Sec. III D), verified against moderate levels of measurement noise (in Supplemental Material Sec. S2 [84]), and demonstrated on experimental electrophysiological data (in Sec. III E), RAD opens the door to studying criticality in noisy real-world settings and thereby connecting our measurements of the world around us to the deeper mechanistic principles that underlie them.

Even though our data-driven methodology for finding noise-robust features can be generically applied to various critical systems, we made a number of simplifying choices that limit how well RAD will perform on arbitrary critical systems. Foremost, we chose to examine a simple normal form [see Sec. II A and Eq. (1)] that describes a broad range of systems, from auditory hair cells [77,78] to financial markets [36], and many others [6,79]. Features that perform well for the radial part of the supercritical Hopf bifurcation, Eq. (1), can be adapted for the full form of the Hopf bifurcation by calculating the radius, and for the pitchfork bifurcation by taking the absolute value; we verified that this centering step allows RAD to translate to the pitchfork bifurcation in Sec. S2 of Supplemental Material [84]. Although numerical simulations were performed only for normal forms, we expect our summary feature to perform well on other systems that exhibit a Hopf or pitchfork bifurcation. However, other normal forms can exhibit fundamentally different changes during bifurcation from the Hopf or pitchfork normal forms. The saddle-node normal form, for instance, occurs when an unstable and a stable equilibrium annihilate one another. In this case, the potential grows more asymmetric as the critical point is approached: Unless the noise is sufficiently weak, the DTC is large, and RAD is unlikely to remain noise robust for saddle-node bifurcations (and other systems with asymmetric normal forms, such as for transcritical bifurcations). Nevertheless, many systems, such as the saddle-node,

transcritical, and subcritical Hopf or pitchfork bifurcations [13]], exhibit critical transitions in which proximity to the critical point corresponds to explosive jumps toward distant attractors. For these systems, time-series features are unable to give a deterministic estimate of the time to catastrophe at high values of η , since crossings of the unstable threshold can be induced by noise well before the critical point [13]. The DTC is still a useful quantity, however, for inferring the likelihood of a critical transition, and given that most potential functions are locally quadratic around stable fixed points, we expect RAD to generically outperform conventional metrics under variable-noise conditions, in particular, for subcritical Hopf and pitchfork bifurcations. Applying our data-driven methodology to find the most noise-robust DTC indicators for new normal forms is a promising avenue for future work, along with using our approach across classes of bifurcations to find a critical indicator that is not only noise robust, but remains consistent over classes of bifurcations [62].

In addition to a simple deterministic component, our model system also bears a highly simplified noise process. Additive dynamical noise, which here is Gaussian, independent, and identically distributed, appears in many systems and is a common modeling assumption. Unlike measurement noise, which is incorporated into the signal after a system has evolved and been measured, dynamical noise is present in the equations of motion for a system, and continually interplays with the deterministic dynamics controlled by the potential function. We verified that RAD is robust to both additive measurement noise, as well as low levels of colored dynamical noise in Sec. S2 of Supplemental Material [84]. Indeed, recent work has analytically studied the problem of estimating the distance to criticality in linear stochastic differential equations with nonstationary colored noise from simulated data [144], developing an algorithm that fits a parametrized power spectrum. In contrast, our data-driven approach demonstrated that simple, noise-robust statistics can be formulated for problems involving analytically intractable nonlinear systems. We also showed that our new metric of criticality has novel and scientifically useful behavior on real-world neural data. Future work may look to extend RAD with multiple time lags to improve its performance in systems with colored noise [144] or subsampled data [99]. Critical systems can also possess non-Gaussian noise [61,145], multiplicative noise, highly nonstationary noise, or even noise in the control parameter itself [6]; we did not study such cases here. Moreover, we did not aim to find features that were insensitive to variation in both the noise amplitude and the sampling period, although we expect this to be a more difficult problem. Regardless, we expect our results will generalize well to systems that (i) are near Hopf or pitchfork bifurcations, (ii) are sufficiently close to a potential minimum, with no extreme perturbations, and (iii) have a noise level small enough for the system to

remain localized around a single potential minimum. For systems outside of this regime, the methodology we used to develop our summary feature—performing a data-driven exploration of candidate features to uncover new theoretical principles—is viable for phase transitions and noise processes of other varieties, and is a promising avenue for future scientific work.

Given the inherent stochasticity and complexity of real-world critical systems, accurately tracking the DTC using statistical properties of time-series data poses a significant challenge, prompting interdisciplinary research efforts spanning several decades. Most work has relied on two substantial assumptions: that the DTC is small enough for the theory of normal forms to apply and that the dynamical noise is negligible or fixed. However, conventional metrics of the DTC are highly sensitive to the more realistic setting of variable, or uncertain, noise. In this work, we used a powerful and thorough data-driven approach that surveys a vast library of time-series features to (i) confirm that conventional metrics are disrupted by a variable-noise amplitude, (ii) uncover unstudied time-series features that are insensitive to the noise amplitude, (iii) scrutinize these noise-robust features to develop new theoretical insight, and (iv) summarize our new understanding of a simple new feature for accurately inferring the DTC in noisy real-world systems. This work thus demonstrated a pragmatic, data-driven way of understanding theoretical systems through simulated data and wide methodological comparison, which can automatically flag promising algorithms to motivate the development of new theory. The result of this process, RAD, introduced here, is a viable measure of DTC for realistic settings of systems corrupted by unknown, and in general variable, noise amplitude. We expect these innovations to enable new applications of dynamical systems thinking to noisy, real-world systems.

The supporting data for the simulations presented in this article are openly available from Figshare [146]. MATLAB code for reproducing our analyses and figures is available at the Criticality repository on GitHub [147], which includes functions for RAD (located in the file `RAD.m`) and `fitSupercriticalHopfRadius_1` (at `potentialDistributions.m`). MATLAB, JULIA, and PYTHON functions for RAD are collected in the RAD repository [91]. Alternatively, RAD is available as the `CR_RAD` function in *hctsa* v1.08 [148] and the `TimeseriesFeatures.jl` JULIA package [149]. RAD is also included in the default set of *hctsa* features as `CR_RAD_1` (centered RAD using lag-1 differences), `CR_RAD_2` (with lag-2 differences), and `CR_RAD_TAU` (using a lag equal to the first zero crossing of the autocorrelation function). For the electrophysiological application presented in Sec. III E, we accessed all Neuropixels Visual Behavior data (manifest v0.5.0) through open software published by the Allen Institute

for Brain Science (the `allensdk` [150]) using the Julia-language wrapper `AllenNeuropixelsBase.jl` [151]. Julia scripts to reproduce our analyses of the neuropixels data are available in the accompanying Criticality repository [147].

ACKNOWLEDGMENTS

We thank Dong-Ping Yang for helpful discussions during the early stages of this project. L. L. G. was supported by the Australian Research Council (ARC) Future Fellowship (FT200100942) and the Rebecca L. Cooper Foundation (PG2019402).

APPENDIX A: INSPECTING HIGH-PERFORMING FEATURES

Here we outline of four *hctsa* features that robustly track the DTC in the presence of uncertain noise (as identified in Sec. III B), which we studied in close to detail to uncover the two algorithmic principles vital for noise-robust inference of the DTC (see Sec. III C). As described in Sec. III B, we manually selected four features from the clustering analysis depicted in Fig. 3, aiming for features that were algorithmically dissimilar and belonged to different clusters. We then studied these four top features in turn to uncover the common algorithmic components that allow for noise-robust inference of the DTC.

1. Change in autocorrelation after discarding extrema

The first feature we investigated was `DN_RemovePoints_max_01_ac1diff`, which measures the change in the autocorrelation of a time series after discarding extreme values. This feature performed well in both the fixed-noise ($\rho_\mu^{\text{fix}} = 0.94$) and variable-noise ($\rho_\mu^{\text{var}} = -0.88$) settings. Acting on an input time series \mathbf{x} , this feature first discards the largest 10% of positive values (while maintaining the temporal ordering of the remaining data points) to produce a thresholded time series \mathbf{x}' . Next, it calculates the lag-1 autocorrelation r_1 for both the original and thresholded time series, outputting the difference of the resulting values as

$$f_{\text{DN}}(\mathbf{x}) = r_1(\mathbf{x}) - r_1(\mathbf{x}'), \quad \mathbf{x}' = \{x_i | x_i < P_{90}(\mathbf{x})\}, \quad (\text{A1})$$

where $P_{90}(\mathbf{x})$ is the 90th percentile of \mathbf{x} . Provided Δt is small compared to the timescale of the deterministic dynamics of a system, it follows from the covariance of the sum of random variables that the lag-1 autocorrelation is approximated by a simple function of only the variance of a time series $\sigma^2(\mathbf{x})$ and the variance of differences $\sigma^2(\Delta\mathbf{x})$. Described by these two quantities, the lag-1 autocorrelation is given by

$$r_1 \approx \frac{2\sigma^2(\mathbf{x}) - \sigma^2(\Delta\mathbf{x})/2}{2\sigma^2(\mathbf{x}) + \sigma^2(\Delta\mathbf{x})/2}. \quad (\text{A2})$$

Thus, `DN_RemovePoints_max_01_ac1diff` relies on (i) properties of the tailedness of the distribution, namely, how $\sigma^2(\mathbf{x})$ differs from $\sigma^2(\mathbf{x}')$, which are (ii) calibrated against the variance of differences $\sigma^2(\Delta\mathbf{x})$. Hence, the algorithmic components of this feature motivated our close study of these two properties, the distribution and the variance of differences, in Sec. III C.

2. Probability of two consecutive high values

A second high-performing feature, in both the fixed-noise ($\rho_\mu^{\text{fix}} = 0.93$) and variable-noise ($\rho_\mu^{\text{var}} = 0.88$) cases, is `SB_MotifTwo_mean_uu`, which counts the proportion of consecutive pairs of time-series values that are both above the mean:

$$f_{\text{SB}} = P(x_t > \bar{x} \cap x_{t+1} > \bar{x}), \quad (\text{A3})$$

where \bar{x} is the mean of \mathbf{x} . Since $\Delta\mathbf{x}$ depends predominantly on the noise amplitude, as described in Sec. III C, the probability of a point crossing the mean over an increment in time depends on (i) its distance from the mean and (ii) the variance of its diffusive motion given by $\sigma^2(\Delta\mathbf{x})$. Hence, by counting the proportion of values that remain above the mean, `SB_MotifTwo_mean_uu` depends on both properties we identified as crucial for robustly inferring the DTC: (i) the proportion of values that are above the mean at the initial time point (related to the distribution) and (ii) the probability of crossing for points above the mean (related to the variance of differences).

3. Change in distribution after moving-average filter

Our third noise-robust estimator for the DTC is `PP_Compare_rav2_kscn_olapint`, which has a high $\rho_\mu^{\text{fix}} = 0.88$ and $\rho_\mu^{\text{var}} = -0.87$. This feature measures the change in the probability density of time-series values (via an overlap integral to a best-fitting Gaussian distribution) after it has been smoothed by a single-period, two-sample moving average:

$$f_{\text{PP}} = \frac{O[p(x; M[\mathbf{x}])]}{O[p(x; \mathbf{x})]}, \quad (\text{A4})$$

where $p(x; \mathbf{x})$ is a kernel-density estimate for the distribution of \mathbf{x} , M is a two-sample moving-average filter, and O is an overlap integral given by

$$O[p(x)] = \sigma(\mathbf{x}) \int p(x) \varphi(\bar{x}, \sigma(\mathbf{x})) dx, \quad (\text{A5})$$

where $\varphi(\mu, \sigma)$ is the probability density of a Gaussian distribution with mean μ and standard deviation σ . The main effect of a moving average filter on the distribution is an increase in the density at medial values, which is a result of peaks and troughs in the time series being truncated and depends on the distribution of values. Moreover, the

magnitude of the effect of truncating the most extreme values also depends on the mean size of the fastest-timescale fluctuations as well as the frequency of extreme peaks, which are both determined by the variance of differences.

4. Comparison of extrema in time-series windows

Finally, we outline `ST_LocalExtrema_1100_meanrat`, which has $\rho_\mu^{\text{fix}} = 0.91$ and $\rho_\mu^{\text{var}} = -0.90$. This feature finds the average of maximum and minimum values computed across nonoverlapping, 100-sample windows of a standardized (z -scored) time series, then returns the ratio of these two values. From another perspective, this feature is a proxy for the relative extent to which the invariant density expands on either side of the mean as η increases. This can be seen by noting that the minimum time-series value in each window is close to zero, such that when the time series is standardized, a value of 0 becomes $-\bar{x}/\sigma(\mathbf{x})$. Therefore, standardization reduces `ST_LocalExtrema_1100_meanrat` to the ratio of the non-normalized average maxima and the mean time-series value:

$$f_{\text{ST}} = \frac{\langle \max[(\mathbf{x}_w - \bar{x})/\sigma(\mathbf{x})] \rangle}{\langle |\min[(\mathbf{x}_w - \bar{x})/\sigma(\mathbf{x})]| \rangle} \approx \frac{\langle \max[\mathbf{x}_w - \bar{x}] \rangle}{\bar{x}}, \quad (\text{A6})$$

where $\langle \max[\mathbf{x}_w] \rangle$ is the mean value of the maximum in each 100-sample window, and $\min[\cdot]$ is the minimum value. As illustrated in Fig. 5, the mean of the distribution (related to the standard deviation) depends on both the DTC and the noise amplitude. The average maximum within a window, however, depends on the rate at which the system diffuses within the short, 100-sample windows. Hence, this final feature, in a similar way to the other top *hctsa* features we have detailed, also compares a distributional property to the variance of differences.

APPENDIX B: CASE STUDY METHODOLOGY

In Sec. III E, we demonstrated how centered RAD applied to electrophysiological data tracks the functional hierarchy of the mouse visual cortex where conventional indicators of criticality and dynamic range failed. Here we provide detailed methods for our case study, including a description of the Allen Visual Behavior Neuropixels recordings, our pipeline for accessing and preprocessing the open dataset, as well as our procedure for evaluating RAD and other metrics on these data.

The Allen Neuropixels Visual Behavior dataset [130] comprises electrophysiological recordings from the mouse visual cortex acquired with Neuropixels probes [152], while the mouse is presented with a series of visual stimuli [153]. A Neuropixels probe houses a checkerboard array of electrodes spaced vertically by $20(\mu\text{m})$ along a linear shank 1 cm in length [152]. Inserted into the mouse brain, each

probe records the electrical activity at hundreds of sites. For the Allen Neuropixels Visual Behavior dataset, six probes were inserted at the retinotopic centers of six areas of the visual cortex in each mouse [153], and raw recordings were filtered into the local field potential (<1000 Hz) and neuronal spikes (0.3–10 kHz). Each of the 75 mice were presented with a series of visual stimuli, including a blank gray screen, Gabor patches, full-field flashes, and an image-change detection task, across two recording sessions. The Neuropixels Visual Behavior dataset having a high recording quality and a large number of probes that cover the majority of areas in the visual cortex provides a valuable opportunity to study the distribution of critical dynamics across the visual hierarchy.

We selected recording sessions that had location-tagged channels in all six of the target visual areas, and mice that displayed no abnormal histology or activity. Session filtering resulted in 39 sessions from 39 mice. From these 39 sessions, we selected all LFP channels located in the cortex, giving between 15 and 24 channels for each visual area, for each mouse. We accessed LFP time series for each channel during the “spontaneous” stimulus, during which the mouse was presented with a static gray screen. The resulting LFP time series were five minutes in length and sampled at a rate of 1250 kHz. We then decimated the signal to a sampling rate of 125 Hz, after bandpass filtering between 1 and 20 Hz to (i) avoid high-frequency down-sampling artifacts as well as (ii) capture the dominant low-frequency activity in the theta (4–8 Hz) band alongside a portion of the higher-frequency $1/f$ activity. Finally, we computed the centered RAD, AC_1, and standard deviation for each cortical channel, for each session, before comparing to the structural hierarchy (as presented in Sec. III E). An exact record of the filtering thresholds and the resultant sessions can be found in the paper/*Criticality.jl* folder of the accompanying code repository [147], alongside scripts and data for reproducing Fig. 6.

[1] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*, 2nd ed. (CRC Press, Boca Raton, 2015).

[2] T. Ma, S. Wang, T. Ma, and S. Wang, *Equilibrium phase transitions in statistical physics*, in *Phase Transition Dynamics* (Springer International Publishing, Cham, 2019), pp. 129–278.

[3] J. M. Carlson, J. T. Chayes, L. Chayes, J. P. Sethna, and D. J. Thouless, *Bethe lattice spin glass: The effects of a ferromagnetic bias and external fields. I. Bifurcation analysis*, *J. Stat. Phys.* **61**, 987 (1990).

[4] A. Mazumdar and G. White, *Review of cosmic phase transitions: Their significance and experimental signatures*, *Rep. Prog. Phys.* **82**, 076901 (2019).

[5] M. A. Muñoz, *Colloquium: Criticality and dynamical scaling in living systems*, *Rev. Mod. Phys.* **90**, 031001 (2018).

[6] F. Freyer, J. A. Roberts, P. Ritter, and M. Breakspear, *A canonical model of multistability and scale-invariance in biological systems*, *PLoS Comput. Biol.* **8**, e1002634 (2012).

[7] L. Cocchi, L. L. Gollo, A. Zalesky, and M. Breakspear, *Criticality in the brain: A synthesis of neurobiology, models and cognition*, *Prog. Neurobiol.* **158**, 132 (2017).

[8] G. Chen and P. Gong, *Computing by modulating spontaneous cortical activity patterns as a mechanism of active visual processing*, *Nat. Commun.* **10**, 4915 (2019).

[9] A. J. Fontenele, N. A. P. de Vasconcelos, T. Feliciano, L. A. A. Aguiar, C. Soares-Cunha, B. Coimbra, L. Dalla Porta, S. Ribeiro, A. J. a. Rodrigues, N. Sousa, P. V. Carelli, and M. Copelli, *Criticality between cortical states*, *Phys. Rev. Lett.* **122**, 208101 (2019).

[10] Y. Xu, A. Schneider, R. Wessel, and K. B. Hengen, *Sleep restores an optimal computational regime in cortical networks*, *Nat. Neurosci.* **27**, 328 (2024).

[11] V. Ajjarapu and B. Lee, *Bifurcation theory and its application to nonlinear dynamical phenomena in an electrical power system*, *IEEE Transactions on Power Systems; IEEE Transactions on Power Electronics* **7**, 424 (1992).

[12] K. N. Srivastava and S. C. Srivastava, *Application of Hopf bifurcation theory for determining critical value of a generator control or load parameter*, *Int. J. Electr. Power Energy Syst.* **17**, 347 (1995).

[13] C. Kuehn, *A mathematical framework for critical transitions: Bifurcations, fast-slow systems and stochastic dynamics*, *Physica (Amsterdam)* **240D**, 1020 (2011).

[14] C. Kuehn, *A mathematical framework for critical transitions: Normal forms, variance and applications*, *J. Nonlinear Sci.* **23**, 457 (2013).

[15] R. D. Peters, M. Le Berre, and Y. Pomeau, *Prediction of catastrophes: An experimental model*, *Phys. Rev. E* **86**, 026207 (2012).

[16] A. J. K. Phillips and P. A. Robinson, *A quantitative model of sleep-wake dynamics based on the physiology of the brainstem ascending arousal system*, *J. Biol. Rhythms* **22**, 167 (2007).

[17] D.-P. Yang, L. McKenzie-Sell, A. Karanjai, and P. A. Robinson, *Wake-sleep transition as a noisy bifurcation*, *Phys. Rev. E* **94**, 022412 (2016).

[18] M. I. Maturana, C. Meisel, K. Dell, P. J. Karoly, W. D’Souza, D. B. Grayden, A. N. Burkitt, P. Jiruska, J. Kudlacek, J. Hlinka, M. J. Cook, L. Kuhlmann, and D. R. Freestone, *Critical slowing down as a biomarker for seizure susceptibility*, *Nat. Commun.* **11**, 2172 (2020).

[19] S. Liu, F. Li, and F. Wan, *Distance to criticality undergoes critical transition before epileptic seizure attacks*, *Brain Research Bulletin* **200**, 110684 (2023).

[20] V. Dakos, M. Scheffer, E. H. van Nes, V. Brovkin, V. Petoukhov, and H. Held, *Slowing down as an early warning signal for abrupt climate change*, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 14308 (2008).

[21] M. Gleiser, E. W. Kolb, and R. Watkins, *Phase transitions with sub-critical bubbles*, *Nucl. Phys.* **B364**, 411 (1991).

[22] L. Reatto, *A complex view of criticality*, *Nat. Phys.* **3**, 594 (2007).

- [23] A. Kianercy, R. Veltri, and K. J. Pienta, *Critical transitions in a game theoretic model of tumour metabolism*, *Interface Focus* **4**, 20140014 (2014).
- [24] C. Trefois, P. M. A. Antony, J. Goncalves, A. Skupin, and R. Balling, *Critical transitions in chronic disease: Transferring concepts from ecology to systems medicine*, *Curr. Opin. Biotechnol.* **34**, 48 (2015).
- [25] A. S. Gsell, U. Scharfenberger, D. Özkundakci, A. Walters, L.-A. Hansson, A. B. G. Janssen, P. Nöges, P. C. Reid, D. E. Schindler, E. Van Donk, V. Dakos, and R. Adrian, *Evaluating early-warning indicators of critical transitions in natural aquatic ecosystems*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, E8089 (2016).
- [26] E. Cotilla-Sanchez, P. D. H. Hines, and C. M. Danforth, *Predicting critical transitions from time series synchrophasor data*, *IEEE Trans. Smart Grid* **3**, 1832 (2012).
- [27] L.-W. Kong, H.-W. Fan, C. Grebogi, and Y.-C. Lai, *Machine learning prediction of critical transition and system collapse*, *Phys. Rev. Res.* **3**, 013090 (2021).
- [28] H. Ren and D. Watts, *Early warning signals for critical transitions in power systems*, *Electric Power Systems Research* **124**, 173 (2015).
- [29] T. M. Lenton, *Early warning of climate tipping points*, *Nat. Clim. Change* **1**, 201 (2011).
- [30] C. W. Lynn and D. S. Bassett, *The physics of brain network structure, function and control*, *Nat. Rev. Phys.* **1**, 318 (2019).
- [31] Y. A. Kuznetsov, *Elements of Applied Bifurcation Theory*, 3rd ed., Applied Mathematical Sciences (Springer, New York, 2004).
- [32] L. Zhang, C. Zhang, and Z. He, *Codimension-one and codimension-two bifurcations of a discrete predator-prey system with strong Allee effect*, *Math. Comput. Simul.* **162**, 155 (2019).
- [33] C. Meisel, A. Klaus, C. Kuehn, and D. Plenz, *Critical slowing down governs the transition to neuron spiking*, *PLoS Comput. Biol.* **11**, e1004097 (2015).
- [34] S. Rau, J. Main, P. Köberle, and G. Wunner, *Pitchfork bifurcations in blood-cell-shaped dipolar Bose-Einstein condensates*, *Phys. Rev. A* **81**, 031605(R) (2010).
- [35] O. V. Ushakov, H.-J. Wünsche, F. Henneberger, I. A. Khovanov, L. Schimansky-Geier, and M. A. Zaks, *Coherence resonance near a Hopf bifurcation*, *Phys. Rev. Lett.* **95**, 123903 (2005).
- [36] Q. Gao and J. Ma, *Chaos and Hopf bifurcation of a finance system*, *Nonlinear Dyn.* **58**, 209 (2009).
- [37] A. J. Hudspeth, *Integrating the active process of hair cells with cochlear function*, *Nat. Rev. Neurosci.* **15**, 600 (2014).
- [38] J. D. Victor, J. D. Drovser, M. M. Conte, and N. D. Schiff, *Mean-field modeling of thalamocortical dynamics and a model-driven approach to EEG analysis*, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 15631 (2011).
- [39] K. M. Aquino, B. Fulcher, S. Oldham, L. Parkes, L. Gollo, G. Deco, and A. Fornito, *On the intersection between data quality and dynamical modelling of large-scale fMRI signals*, *NeuroImage* **256**, 119051 (2022).
- [40] J. D. Crawford, *Introduction to bifurcation theory*, *Rev. Mod. Phys.* **63**, 991 (1991).
- [41] J. O'Byrne and K. Jerbi, *How critical is brain criticality?*, *Trends Neurosci.* **45**, 820 (2022).
- [42] M. Scheffer, S. R. Carpenter, T. M. Lenton, J. Bascompte, W. Brock, V. Dakos, J. Van De Koppel, I. A. Van De Leemput, S. A. Levin, E. H. Van Nes, M. Pascual, and J. Vandermeer, *Anticipating critical transitions*, *Science* **338**, 344 (2012).
- [43] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara, *Early-warning signals for critical transitions*, *Nature (London)* **461**, 53 (2009).
- [44] Y. Hart, M. Vaziri-Pashkam, and L. Mahadevan, *Early warning signals in motion inference*, *PLoS Comput. Biol.* **16**, e1007821 (2020).
- [45] C. Wissel, *A universal law of the characteristic return time near thresholds*, *Oecologia* **65**, 101 (1984).
- [46] E. Negahbani, D. A. Steyn-Ross, M. L. Steyn-Ross, M. T. Wilson, and J. W. Sleigh, *Noise-induced precursors of state transitions in the stochastic Wilson-Cowan model*, *J. Math. Neurosci.* **5**, 9 (2015).
- [47] S. M. O'Regan and D. L. Burton, *How stochasticity influences leading indicators of critical transitions*, *Bull. Math. Biol.* **80**, 1630 (2018).
- [48] V. Dakos, E. H. van Nes, P. D'Odorico, and M. Scheffer, *Robustness of variance and autocorrelation as indicators of critical slowing down*, *Ecology* **93**, 264 (2012).
- [49] S. R. Carpenter, J. J. Cole, M. L. Pace, R. Batt, W. A. Brock, T. Cline, J. Coloso, J. R. Hodgson, J. F. Kitchell, D. A. Seekell, L. Smith, and B. Weidel, *Early warnings of regime shifts: A whole-ecosystem experiment*, *Science* **332**, 1079 (2011).
- [50] V. Dakos, E. H. van Nes, R. Donangelo, H. Fort, and M. Scheffer, *Spatial correlation as leading indicator of catastrophic shifts*, *Theor. Ecol.* **3**, 163 (2009).
- [51] E. Weinans, R. Quax, E. H. van Nes, and I. A. v. d. Leemput, *Evaluating the performance of multivariate indicators of resilience loss*, *Sci. Rep.* **11**, 9148 (2021).
- [52] J. Liu, Y. Tao, R. Lan, J. Zhong, R. Liu, and P. Chen, *Identifying the critical state of cancers by single-sample Markov flow entropy*, *PeerJ Comput. Sci.* **11**, e15695 (2023).
- [53] L. Arnold, *Random Dynamical Systems*, 1st ed., Monographs in Mathematics (Springer, New York, 1998).
- [54] M. Siefert, A. Kittel, R. Friedrich, and J. Peinke, *On a quantitative method to analyze dynamical and measurement noise*, *Europhys. Lett.* **61**, 466 (2003).
- [55] T. Sase, J. P. Ramírez, K. Kitajo, K. Aihara, and Y. Hirata, *Estimating the level of dynamical noise in time series by using fractal dimensions*, *Phys. Lett. A* **380**, 1151 (2016).
- [56] H. Crauel and F. Flandoli, *Additive noise destroys a pitchfork bifurcation*, *J. Dyn. Differ. Equ.* **10**, 259 (1998).
- [57] M. Callaway, T. S. Doan, J. S. W. Lamb, and M. Rasmussen, *The dichotomy spectrum for random dynamical systems and pitchfork bifurcations with additive noise*, *Ann. Inst. Henri Poincaré B* **53**, 1548 (2017).
- [58] C. Meunier and A. D. Verga, *Noise and bifurcations*, *J. Stat. Phys.* **50**, 345 (1988).
- [59] J. M. T. Thompson and J. Sieber, *Predicting climate tipping as a noisy bifurcation: A review*, *Int. J. Bifurcation Chaos Appl. Sci. Eng.* **21**, 399 (2011).

- [60] T. Wilkat, T. Rings, and K. Lehnertz, *No evidence for critical slowing down prior to human epileptic seizures*, *Chaos* **29**, 091104 (2019).
- [61] L. Fronzoni, R. Mannella, P. V. E. McClintock, and F. Moss, *Postponement of Hopf bifurcations by multiplicative colored noise*, *Phys. Rev. A* **36**, 834 (1987).
- [62] F. Grziwotz, C.-W. Chang, V. Dakos, E. H. van Nes, M. Schwarzländer, O. Kamps, M. Heßler, I. T. Tokuda, A. Telschow, and C.-h. Hsieh, *Anticipating the occurrence and type of critical transitions*, *Sci. Adv.* **9**, eabq4558 (2023).
- [63] O. van der Groen, J. B. Mattingley, and N. Wenderoth, *Altering brain dynamics with transcranial random noise stimulation*, *Sci. Rep.* **9**, 4029 (2019).
- [64] D.-P. Yang and P. A. Robinson, *Critical dynamics of Hopf bifurcations in the corticothalamic system: Transitions from normal arousal states to epileptic seizures*, *Phys. Rev. E* **95**, 042410 (2017).
- [65] J. M. T. Thompson and J. Sieber, *Climate tipping as a noisy bifurcation: A predictive technique*, *IMA J. Appl. Math.* **76**, 27 (2010).
- [66] J. F. Mejjias and A. Longtin, *Differential effects of excitatory and inhibitory heterogeneity on the gain and asynchronous state of sparse cortical networks*, *Front. Comput. Neurosci.* **8**, 107 (2014).
- [67] B. D. Fulcher, M. A. Little, and N. S. Jones, *Highly comparative time-series analysis: The empirical structure of time series and their methods*, *J. R. Soc. Interface* **10**, 20130048 (2013).
- [68] B. D. Fulcher and N. S. Jones, *hctsa: A computational framework for automated time-series phenotyping using massive feature extraction*, *Cell Syst.* **5**, 527 (2017).
- [69] B. D. Fulcher, *Feature-based time-series analysis*, in *Feature Engineering for Machine Learning and Data Analytics*, 1st ed., edited by G. Dong and H. Liu (CRC Press, Boca Raton, 2018), Chap. 4, pp. 87–109.
- [70] C. W. Tan, C. Bergmeir, F. Petitjean, and G. I. Webb, *Time series extrinsic regression: Predicting numeric values from time series data*, *Data Min. Knowl. Discov.* **35**, 1032 (2021).
- [71] S. J. Kiebel, J. Daunizeau, and K. J. Friston, *A hierarchy of time-scales and the brain*, *PLoS Comput. Biol.* **4**, e1000209 (2008).
- [72] J. D. Murray, A. Bernacchia, D. J. Freedman, R. Romo, J. D. Wallis, X. Cai, C. Padoa-Schioppa, T. Pasternak, H. Seo, D. Lee, and X.-J. Wang, *A hierarchy of intrinsic timescales across primate cortex*, *Nat. Neurosci.* **17**, 1661 (2014).
- [73] L. L. Gollo, A. Zalesky, R. M. Hutchison, M. van den Heuvel, and M. Breakspear, *Dwelling quietly in the rich club: Brain network determinants of slow cortical fluctuations*, *Phil. Trans. R. Soc. B* **370**, 20140165 (2015).
- [74] R. Chaudhuri, K. Knoblauch, M.-A. Gariel, H. Kennedy, and X.-J. Wang, *A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex*, *Neuron* **88**, 419 (2015).
- [75] R. V. Raut, A. Z. Snyder, and M. E. Raichle, *Hierarchical dynamics as a macroscopic organizing principle of the human brain*, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 20890 (2020).
- [76] J. Liang, Z. Yang, and C. Zhou, *Excitation–inhibition balance, neural criticality, and activities in neuronal circuits*, *Neuroscientist* (2024), [10.1177/10738584231221766](https://doi.org/10.1177/10738584231221766).
- [77] M. Ospeck, V. M. Eguíluz, and M. O. Magnasco, *Evidence of a Hopf bifurcation in frog hair cells*, *Biophys. J.* **80**, 2597 (2001).
- [78] D. Ó. Maoiléidigh, E. M. Nicola, and A. J. Hudspeth, *The diverse effects of mechanical loading on active hair bundles*, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 1943 (2012).
- [79] J. E. Marsden, M. McCracken, J. E. Marsden, and M. McCracken, *Examples, in The Hopf Bifurcation and Its Applications* (Springer, New York, 1976), pp. 136–150.
- [80] G. Deco, M. L. Kringelbach, V. K. Jirsa, and P. Ritter, *The dynamics of resting fluctuations in the brain: Metastability and its dynamical cortical core*, *Sci. Rep.* **7**, 3095 (2017).
- [81] M. Schumaker, *Center manifold reduction and normal form transformations in systems with additive noise*, *Phys. Lett. A* **122**, 317 (1987).
- [82] J. Ferrell, *Bistability, bifurcations, and Waddington’s epigenetic landscape*, *Curr. Biol.* **22**, R458 (2012).
- [83] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations* (Springer, Berlin, 1992).
- [84] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.14.031021> for a joint analysis of feature performance under fixed and variable noise, as well as validation of RAD on additional systems.
- [85] C. M. Thomas and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (John Wiley & Sons, Nashville, 2006).
- [86] S. R. Carpenter and W. A. Brock, *Rising variance: A leading indicator of ecological transition*, *Ecol. Lett.* **9**, 311 (2006).
- [87] Y.-C. Chen, *A tutorial on kernel density estimation and recent advances*, *Biostat. Epidemiol.* **1**, 161 (2017).
- [88] J. D. Hamilton, *Time Series Analysis* (Princeton University Press, Princeton, NJ, 1994).
- [89] R. F. Bass, *Stochastic Processes*, Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press, Cambridge, England, 2011).
- [90] F. Takens, *Detecting strange attractors in turbulence*, in *Dynamical Systems and Turbulence, Warwick 1980*, Lecture Notes in Mathematics Vol. 898, edited by D. Rand and L.-S. Young (Springer, Berlin, 1981), Chap. 21, pp. 366–381.
- [91] B. Harris, *DynamicsAndNeuralSystems/RAD: v0.1.0* (2024), [10.5281/zenodo.12660860](https://doi.org/10.5281/zenodo.12660860).
- [92] O. Kinouchi and M. Copelli, *Optimal dynamical range of excitable networks at criticality*, *Nat. Phys.* **2**, 348 (2006).
- [93] S. H. Gautam, T. T. Hoang, K. McClanahan, S. K. Grady, and W. L. Shew, *Maximizing sensory dynamic range by tuning the cortical state to criticality*, *PLoS Comput. Biol.* **11**, e1004576 (2015).
- [94] W. L. Shew, H. Yang, T. Petermann, R. Roy, and D. Plenz, *Neuronal avalanches imply maximum dynamic range in cortical networks at criticality*, *J. Neurosci.* **29**, 15595 (2009).
- [95] L. L. Gollo, *Coexistence of critical sensitivity and sub-critical specificity can yield optimal population coding*, *J. R. Soc. Interface* **14**, 20170207 (2017).
- [96] W. L. Shew and D. Plenz, *The functional benefits of criticality in the cortex*, *Neurosci.* **19**, 88 (2012).

- [97] L. L. Gollo, O. Kinouchi, and M. Copelli, *Single-neuron criticality optimizes analog dendritic computation*, *Sci. Rep.* **3**, 3222 (2013).
- [98] A. Gal and S. Marom, *Self-organized criticality in single-neuron excitability*, *Phys. Rev. E* **88**, 062717 (2013).
- [99] J. Wilting and V. Priesemann, *Inferring collective dynamical states from widely unobserved systems*, *Nat. Commun.* **9**, 2325 (2018).
- [100] Z. Ma, G. G. Turrigiano, R. Wessel, and K. B. Hengen, *Cortical circuit dynamics are homeostatically tuned to criticality in-vivo*, *Neuron* **104**, 655 (2019).
- [101] D. Dahmen, S. Grün, M. Diesmann, and M. Helias, *Second type of criticality in the brain uncovers rich multiple-neuron dynamics*, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 13051 (2019).
- [102] B. Munn, N. Zeater, A. N. Pietersen, S. G. Solomon, S. K. Cheong, P. R. Martin, and P. Gong, *Fractal spike dynamics and neuronal coupling in the primate visual system*, *J. Physiol.* **598**, 1551 (2020).
- [103] N. Lotfi, T. Feliciano, L. A. A. Aguiar, Thais Priscila Lima Silva, T. T. A. Carvalho, O. A. Rosso, M. Copelli, F. S. Matias, and P. V. Carelli, *Statistical complexity is maximized close to criticality in cortical dynamics*, *Phys. Rev. E* **103**, 012415 (2021).
- [104] G. B. Morales, S. di Santo, and M. A. Muñoz, *Quasiuniversal scaling in mouse-brain neuronal activity stems from edge-of-instability critical dynamics*, *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2208998120 (2023).
- [105] S. A. Jones, J. H. Barfield, V. K. Norman, and W. L. Shew, *Scale-free behavioral dynamics directly linked with scale-free cortical dynamics*, *eLife* **12**, e79950 (2023).
- [106] D. M. Castro, T. Feliciano, N. A. P. de Vasconcelos, C. Soares-Cunha, B. Coimbra, A. J. a. Rodrigues, P. V. Carelli, and M. Copelli, *In and out of criticality? State-dependent scaling in the rat visual cortex*, *PRX Life* **2**, 023008 (2024).
- [107] W. L. Shew, W. P. Clawson, J. Pobst, Y. Karimipناه, N. C. Wright, and R. Wessel, *Adaptation to sensory input tunes visual cortex to criticality*, *Nat. Phys.* **11**, 659 (2015).
- [108] G. Hahn, A. Ponce-Alvarez, C. Monier, G. Benvenuti, A. Kumar, F. Chavane, G. Deco, and Y. Frégnac, *Spontaneous cortical activity is transiently poised close to criticality*, *PLoS Comput. Biol.* **13**, e1005543 (2017).
- [109] B. Mariani, G. Nicoletti, M. Bisio, M. Maschietto, S. Vassanelli, and S. Suweis, *Disentangling the critical signatures of neural activity*, *Sci. Rep.* **12**, 10770 (2022).
- [110] M. Fuscà, F. Siebenhühner, S. H. Wang, V. Myrov, G. Arnulfo, L. Nobili, J. M. Palva, and S. Palva, *Brain criticality predicts individual levels of inter-areal synchronization in human electrophysiological data*, *Nat. Commun.* **14**, 4736 (2023).
- [111] E. Tagliazucchi, P. Balenzuela, D. Fraiman, and D. R. Chialvo, *Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis*, *Front. Oral Physiol.* **3**, 15 (2012).
- [112] J. O'Byrne and K. Jerbi, *How critical is brain criticality?*, *Trends Neurosci.* **45**, 820 (2022).
- [113] V. Priesemann, M. Wibral, M. Valderrama, R. Pröpper, M. Le Van Quyen, T. Geisel, J. Triesch, D. Nikolić, and M. H. Munk, *Spike avalanches in vivo suggest a driven, slightly subcritical brain state*, *Front. Syst. Neurosci.* **8**, 108 (2014).
- [114] J. Wilting and V. Priesemann, *25 years of criticality in neuroscience-established results, open controversies, novel concepts*, *Curr. Opin. Neurobiol.* **58**, 105 (2019).
- [115] N. Tomen, D. Rotermund, and U. Ernst, *Marginally subcritical dynamics explain enhanced stimulus discriminability under attention*, *Front. Syst. Neurosci.* **8**, 151 (2014).
- [116] E. D. Fagerholm, R. Lorenz, G. Scott, M. Dinov, P. J. Hellyer, N. Mirzaei, C. Leeson, D. W. Carmichael, D. J. Sharp, W. L. Shew *et al.*, *Cascades and cognitive state: Focused attention incurs subcritical dynamics*, *J. Neurosci.* **35**, 4626 (2015).
- [117] C. C. Hilgetag and A. Goulas, *'Hierarchy' in the organization of brain networks*, *Phil. Trans. R. Soc. B* **375**, 20190319 (2020).
- [118] R. D. D'Souza, Q. Wang, W. Ji, A. M. Meier, H. Kennedy, K. Knoblauch, and A. Burkhalter, *Hierarchical and non-hierarchical features of the mouse visual cortical network*, *Nat. Commun.* **13**, 503 (2022).
- [119] L. L. Gollo, *Exploring atypical timescales in the brain*, *eLife* **8**, e45089 (2019).
- [120] R. D. D'Souza, A. M. Meier, P. Bista, Q. Wang, and A. Burkhalter, *Recruitment of inhibition and excitation across mouse visual cortex depends on the hierarchy of inter-connecting areas*, *eLife* **5**, e19332 (2016).
- [121] B. D. Fulcher, J. D. Murray, V. Zerbi, and X.-J. Wang, *Multimodal gradients across mouse cortex*, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 4689 (2019).
- [122] L. Cocchi, M. V. Sale, L. L. Gollo, P. T. Bell, V. T. Nguyen, A. Zalesky, M. Breakspear, and J. B. Mattingley, *A hierarchy of timescales explains distinct effects of local inhibition of primary visual cortex and frontal eye fields*, *eLife* **5**, e15252 (2016).
- [123] L. L. Glickfeld and S. R. Olsen, *Higher-order areas of the mouse visual cortex*, *Annu. Rev. Vis. Sci.* **3**, 251 (2017).
- [124] Y. Tian, Z. Tan, H. Hou, G. Li, A. Cheng, Y. Qiu, K. Weng, C. Chen, and P. Sun, *Theoretical foundations of studying criticality in the brain*, *Netw. Neurosci.* **6**, 1148 (2022).
- [125] H. Young, B. Belbut, M. Baeta, and L. Petreanu, *Laminar-specific cortico-cortical loops in mouse visual cortex*, *eLife* **10**, e59551 (2021).
- [126] Y.-t. Wu, H. C. Bennett, U. Chon, D. J. Vanselow, Q. Zhang, R. Muñoz-Castañeda, K. C. Cheng, P. Osten, P. J. Drew, and Y. Kim, *Quantitative relationship between cerebrovascular network and neuronal cell types in mice*, *Cell Rep.* **39**, 110978 (2022).
- [127] G. Shafei, B. D. Fulcher, B. Voytek, T. D. Satterthwaite, S. Baillet, and B. Misić, *Neurophysiological signatures of cortical micro-architecture*, *Nat. Commun.* **14**, 6000 (2023).
- [128] J. H. Siegle *et al.*, *Survey of spiking in the mouse visual system reveals functional hierarchy*, *Nature (London)* **592**, 86 (2021).
- [129] J. A. Harris *et al.*, *Hierarchical organization of cortical and thalamic connectivity*, *Nature (London)* **575**, 195 (2019).
- [130] Allen Institute MindScope Program, *Allen Institute—visual behavior—neuropixels, Version 0.240702.1725*, DANDI archive (2024), [10.48324/dandi.000713/0.240702.1725](https://doi.org/10.48324/dandi.000713/0.240702.1725).
- [131] B. Lindner, *Effects of noise in excitable systems*, *Phys. Rep.* **392**, 321 (2004).

- [132] D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G.-Z. Yang, *XAI—explainable artificial intelligence*, *Sci. Robot.* **4**, eaay7120 (2019).
- [133] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, *Explainable AI: A review of machine learning interpretability methods*, *Entropy* **23**, 18 (2020).
- [134] A. Davies, P. Veličković, L. Buesing, S. Blackwell, D. Zheng, N. Tomašev, R. Tanburn, P. Battaglia, C. Blundell, A. Juhász, M. Lackenby, G. Williamson, D. Hassabis, and P. Kohli, *Advancing mathematics by guiding human intuition with AI*, *Nature (London)* **600**, 70 (2021).
- [135] C. Tu, P. D’Odorico, and S. Suweis, *Critical slowing down associated with critical transition and risk of collapse in crypto-currency*, *R. Soc. Open Sci.* **7**, 191450 (2022).
- [136] X.-Q. Xie, W.-P. He, B. Gu, Y. Mei, and J. Wang, *The robustness of the skewness as an early warning signal for abrupt climate change*, *International Journal of Climatology* **39**, 5672 (2019).
- [137] V. Guttal and C. Jayaprakash, *Changing skewness: An early warning signal of regime shifts in ecosystems*, *Ecol. Lett.* **11**, 450 (2008).
- [138] I. A. van de Leemput, M. Wichers, A. O. J. Cramer, D. Borsboom, F. Tuerlinckx, P. Kuppens, E. H. van Nes, W. Viechtbauer, E. J. Giltay, S. H. Aggen, C. Derom, N. Jacobs, K. S. Kendler, H. L. J. van der Maas, M. C. Neale, F. Peeters, E. Thiery, P. Zachar, and M. Scheffer, *Critical slowing down as early warning for the onset and termination of depression*, *Proc. Natl. Acad. Sci. U.S.A.* **111**, 87 (2014).
- [139] S. Deb, S. Sidheekh, C. F. Clements, N. C. Krishnan, and P. S. Dutta, *Machine learning methods trained on simple models can predict critical transitions in complex natural systems*, *R. Soc. Open Sci.* **9**, 211475 (2022).
- [140] J. Lim and B. I. Epureanu, *Forecasting a class of bifurcations: Theory and experiment*, *Phys. Rev. E* **83**, 016203 (2011).
- [141] M. Markicevic, I. Savvateev, C. Grimm, and V. Zerbi, *Emerging imaging methods to study whole-brain function in rodent models*, *Transl. Psychiatry* **11**, 457 (2021).
- [142] V. Zimmern, *Why brain criticality is clinically relevant: A scoping review*, *Front. Neural Circuits* **14**, 54 (2020).
- [143] S. Trakoshis *et al.*, *Intrinsic excitation-inhibition imbalance affects medial prefrontal cortex differently in autistic men versus women*, *eLife* **9**, e55684 (2020).
- [144] A. Morr and N. Boers, *Detection of approaching critical transitions in natural systems driven by red noise*, *Phys. Rev. X* **14**, 021037 (2024).
- [145] T. Kaur and P. S. Dutta, *Effects of noise correlation and imperfect data sampling on indicators of critical slowing down*, *Theor. Ecol.* **15**, 129 (2022).
- [146] B. Harris, *hctsa results for the radial component of the supercritical Hopf normal form*, (2024), https://figshare.com/articles/dataset_i_hctsa_i_results_for_the_radial_component_of_the_supercritical_Hopf_normal_form/23995104.
- [147] B. Harris, *brendanjohnharris/Criticality: v0.2.1*, Zenodo (2024), [10.5281/zenodo.10968712](https://zenodo.org/record/10968712).
- [148] B. Fulcher, O. Cliff, B. Harris, Philiphorst, S. Sethi, C. H. Lubba, I. Alam, lukas, K. P. Vysyaraju, and J. McCormac, VP007-py, XavierFPMorris, and K. Shiina, *benfulcher/hctsa: v0.98*, Zenodo (2018), [10.5281/zenodo.3927083](https://zenodo.org/record/3927083).
- [149] B. Harris, *brendanjohnharris/TimeSeriesFeatures.jl: v0.4.0*, Zenodo, (2023), [10.5281/zenodo.10551308](https://zenodo.org/record/10551308).
- [150] <https://github.com/AllenInstitute/AllenSDK>.
- [151] B. Harris, *brendanjohnharris/AllenNeuropixelsBase.jl: v0.1.0*, Zenodo (2024), [10.5281/zenodo.10584984](https://zenodo.org/record/10584984).
- [152] J. J. Jun *et al.*, *Fully integrated silicon probes for high-density recording of neural activity*, *Nature (London)* **551**, 232 (2017).
- [153] Allen Institute for Brain Science, *Allen Brain Observatory: Visual Behavior Neuropixels Technical Whitepaper* (Allen Institute for Brain Science, Seattle, Washington State, 2022).