CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	Classifying seizure generation mechanisms: A critical transitions framework
Authors	Flynn, Andrew;McCafferty, Cian;Lehnertz, Klaus;David, François;Crunelli, Vincenzo;Marnane, William P.;Wieczorek, Sebastian
Publication date	2025-11-25
Original Citation	Flynn, A, McCafferty, C, Lehnertz, K, David, F, Crunelli, V, Marnane, W P & Wieczorek, S 2025 'Classifying seizure generation mechanisms : A critical transitions framework' arXiv, pp. 1-42. https://doi.org/10.48550/arXiv.2511.20522
Type of publication	Article (preprint)
Link to publisher's version	10.48550/arXiv.2511.20522
Rights	https://creativecommons.org/licenses/by/4.0/
Download date	2026-10-25 08:28:09
Item downloaded from	https://hdl.handle.net/10468/18454

UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Classifying seizure generation mechanisms: A critical transitions framework.

New insights from combining brain voltage recordings, models of critical transitions and machine learning.

Andrew Flynn[1,2,*], Cian McCafferty[3], Klaus Lehnertz[4,5,6], François David[7], Vincenzo Crunelli[8], William P. Marnane[2,9], and Sebastian Wieczorek[1]

[1]School of Mathematical Sciences, University College Cork, T12 XF62 Cork, Ireland.
[2]INFANT Research Centre, University College Cork, T12 DC4A Cork, Ireland.
[3]Department of Anatomy & Neuroscience, University College Cork, Cork, Ireland.
[4]Department of Epileptology, University of Bonn Medical Centre, 53127 Bonn, Germany.
[5]Helmholtz-Institute for Radiation and Nuclear Physics, University of Bonn, 53115 Bonn, Germany.
[6]Interdisciplinary Center for Complex Systems, University of Bonn, 53175 Bonn, Germany.
[7]Center for Interdisciplinary Research in Biology, Collège de France, 75005 Paris, France
[8]Department of Pharmacology and Neuroscience, Faculty of Medicine, University of Lisbon, Lisbon, Portugal
[9]Electrical and Electronic Engineering, University College Cork, T12 YF78 Cork, Ireland.
[*]andrewflynn@ucc.ie

arXiv:2511.20522v1 [math.DS] 25 Nov 2025

ABSTRACT

Understanding how the brain switches from normal activity to an epileptic seizure is essential for improving seizure therapy, yet the underlying mechanisms remain largely unknown. In particular, seizure onset can be described as a critical transition (CT), but there is no consensus on whether (i) bifurcation-induced, (ii) noise-induced, or (iii) bifurcation/noise-induced CTs are responsible. To clarify this, we develop a versatile CT-classification framework that can be applied to seizures in both animals and humans. First, we identify a canonical mathematical model which displays CTs that closely resemble voltage recordings of real seizures and can be of the three types mentioned above. We then identify distinctive properties of each CT-type in the model's output and use them to train a machine learning CT-type classifier. Finally, we apply the model-trained classifier to voltage recordings from epileptic rodents. We find that the largest proportion of analysed seizures are classified as noise-induced CTs. This challenges the conventional view that seizures are predominantly bifurcation-induced and could inform new therapeutic strategies for seizures.

Introduction

Epileptic seizures are states of high-amplitude, synchronous electrical activity in the brain with defined onset and offset[1]. Seizures have a wide range of associated symptoms but all constitute a large change in brain state and have a commensurate impact on the life of a person with epilepsy[2]. Why the brain suddenly changes from normal activity to an epileptic seizure is unknown[3]. Understanding this process may lead to the more optimal use of existing therapeutic interventions and advance the development of new ones that either prevent onset of seizure or accelerate its end. A sudden and large change in the state of a complex system exactly fits Ashwin *et al.*'s[4] definition of a critical transition (CT). The mathematical theory of CTs is perhaps best known for describing and foreseeing sudden and large changes in climate and ecological systems[5,6]. It has provided a better understanding of key tipping elements in the Earth system and their impacts on our lives, and has also informed global policymaking[7,8] Inspired by its impact on the environmental sciences, we apply the theory of CTs to analyse seizure activity in the brain. Specifically, we model the onset of seizure activity as a CT between two different states of the brain, the non-seizure state (NS state) and the seizure state (S state).

Despite many previous efforts, no consensus has emerged on what types of CT are predominantly responsible for the onset of seizures in the brain. Scheffer *et al.*[9] reported that, in an example of an electroencephalography (EEG) time series with seizure activity taken from McSharry *et al.*[10], the variance of the time series increased prior to seizure onset. From the theory of CTs, this increase in variance is a signature of 'critical slowing down' (CSD), a phenomenon that occurs prior to a bifurcation-induced CT, thus indicating that the seizure occurred due to a bifurcation taking place in the brain. Meisel and Kuehn[11] later analysed an EEG time series from eight adults with epilepsy and found similar increases in variance before only two out of eight seizures. They also showed that such increases occur prior to a subcritical Hopf bifurcation in a mathematical model. Milanowski and Suffczynski[12] expanded on this approach by analysing hundreds of EEG time series of various seizure types in adult humans. They restricted their analysis to CTs due to crossing different Hopf bifurcations and concluded that seizures start without common signatures of such CTs. Jirsa *et al.*[13] used a different modelling approach, based on noise-free dynamic bifurcations, to argue that seizures occur due to slow passage through a saddle-node bifurcation. In more recent times, further opposing reports emerged based on large-scale investigations of EEG time series: Wilkat *et al.*[14] found no evidence of CSD before seizures, whereas Maturana *et al.*[15] did. From a CT point of view, if a bifurcation does not appear to be responsible for the onset of the seizure, then another likely candidate is a noise-induced CT, which have long been theorised to play a role in seizure onset[16]. However, very little is known about the typical signatures of noise-induced CTs in a time series, nor how these signatures manifest themselves in voltage recordings of seizure activity[17].

In this paper, we approach the question of 'how are seizures generated?' from a broader perspective. Specifically, we go beyond the realm of bifurcation-induced CTs and hypothesise that the mechanism of seizure generation is a CT from the NS to S state that can be: noise-induced, bifurcation-induced, or a combination of both. In other words, we consider these three CT-types as potential seizure generation mechanisms. To test this hypothesis, we introduce a new CT-based framework, summarised in four steps in Table 1, to analyse voltage recordings of real seizure activity. This framework consists of an algorithm to detect CTs between NS and S states, and a machine learning classifier to distinguish between the different CT-types. This enables us to test whether, at a given moment in time near a CT to the S state, the brain: (i) operates within a multistable regime whereby noise alone in the form of exogenous or endogenous disturbances can cause it to cross some threshold and trigger a seizure, (ii) slowly drifts towards and then past a bifurcation point that generates the seizure, or (iii) experiences a combination of both (i) and (ii). Seen in this light, our framework can be used not only to address a fundamental question regarding seizure generation mechanisms, but also to inform future seizure prevention and treatment strategies - once we know the CT types that are responsible, we can relate them to different types of neural activity, and adapt therapeutic strategies accordingly.

The structure of the Results section reflects our four-step framework from Table 1. In the first step, we identify a set of two differential equations (an augmented Bautin bifurcation normal form) as a canonical mathematical model capable of generating the three CT-types listed above. We then show that the parameters of our mathematical model can be tuned so that its output resembles real seizure activity in voltage recordings taken from inside the brain of epileptic rodents. More specifically, we use local field potential (LFP) measurements of neural activity in Genetic Absence Epilepsy Rats from Strasbourg (GAERS) that were obtained and annotated by McCafferty *et al.*[18]. We refer to these LFP measurements as the 'voltage recordings' throughout the paper. We choose to work with GAERS because their absence seizures are fully generalised and morphologically distinctive in EEG[19] compared to other seizures which may be focal or have variable EEG signatures[20]. In other words, GAERS seizures are both stable in amplitude and frequent in occurrence[21], making them ideal for investigating seizure onset dynamics and for constructing our framework. We see this as a proof of concept and an important step towards analysing more complex human seizures. In the second step, we construct an algorithm based on the concept of a non-ideal relay to detect CTs between the NS and S states in noisy time series. We test the algorithm on our mathematical model and, crucially, show that it detects CTs between NS and S states in the actual voltage recordings from different GAERS in excellent agreement with expert annotations. We then identify several distinctive properties in time series of the mathematical model's output that are associated with each CT-type. This is one of the most important components of our study. It allows us, in the third step, to train a support vector

Step 1: Build mathematical model of CTs with real seizure characteristics	Construct a canonical mathematical model that displays CTs from its NS to S state that have characteristics of real seizure activity and can be of three types: (i) bifurcation-induced, (ii) noise-induced, (iii) bifurcation/noise-induced.
Step 2: Detect model-generated CTs and identify distinctive properties of each CT-type	Use the mathematical model from Step 1 to generate many examples of CTs from its NS to S state, where the CT-type is known. Develop a CT detection algorithm to detect CTs from the NS to S state in the model's output and group them by their type. For detected CTs of each type, identify their characteristic time series properties.
Step 3: Build CT-type classifier	Train a machine learning classifier using the characteristic time series properties identified in the model's output for each CT-type in Step 2. Optimise the classifier by maximising its ability to correctly classify the types of CT from the NS to S state in the unseen model's output.
Step 4: Detect real seizures and classify their CT-type.	Use the CT detection algorithm from Step 2 to detect CTs from the NS to S state in the voltage recordings, where the CT-type is not known. For each detected CT, calculate the characteristic time series properties identified in Step 2 and use these properties as an input for the model-trained classifier from Step 3 to classify the CT-type.

Table 1. A summary of our proposed four-step framework for classifying CTs that are responsible for the onset of seizures in the voltage recordings of brain activity as: (i) bifurcation-induced, (ii) noise-induced or (iii) bifurcation/noise-induced.

machine classifier (SVM classifier) to distinguish between the three CT-types from the NS to S state in a time series, without any prior knowledge of CT-types. In the fourth and final step, we use the model-trained SVM classifier to identify the types of CT that are responsible for real seizure onsets in the voltage recordings from different GAERS.

Our findings demonstrate that: (i) seizures are generated by different types of CT and (ii) noise-induced CTs appear to be the dominant mechanism responsible for generating the seizures we analysed. These results align with recent experiments on seizure termination in GAERS[22], challenge the conventional view that seizures occur predominantly due to bifurcation-induced CTs, and highlight the need to account for noise-induced CTs when developing seizure therapies. Furthermore, our framework is versatile in the following ways: (i) it can use other, possibly more realistic, mathematical models of seizure activity to train the SVM, (ii) it can classify CT-types responsible for other types of seizures, including those in humans, and (iii) it can also be applied to similar instabilities found in complex systems beyond the brain.

Results

The main outcome of this paper is a CT-type classification framework consisting of:

- A CT-based formulation of how the brain transitions between non-seizure and seizure states.
- An algorithm that detects seizure onset and offset times in voltage recordings of real seizure activity.
- A CT-type classifier based on explainable machine learning tools that classifies seizure generation mechanisms in voltage recordings of real seizure activity.

The main result is as follows: Comparing animal seizure data with a mathematical model within our classification framework, we found that seizures are predominantly generated by noise-induced CTs. This aligns with recent experiments, challenges recent trends that refer to seizures as bifurcation-induced CTs, and could inform future seizure therapies.

Step 1: A mathematical model with real seizure characteristics and three CT types

In this section, we describe three important characteristics of real seizure activity in the voltage recordings from GAERS. We then introduce a canonical mathematical model that mimics these three characteristics and also show the three CT types.

Characteristics of voltage recordings of real seizure activity

Before following the 4 steps in Table 1, we briefly describe three important characteristics of real seizure activity in the voltage recordings that we tune our mathematical model to mimic. These characteristics allow us to (i) use the same algorithm to detect

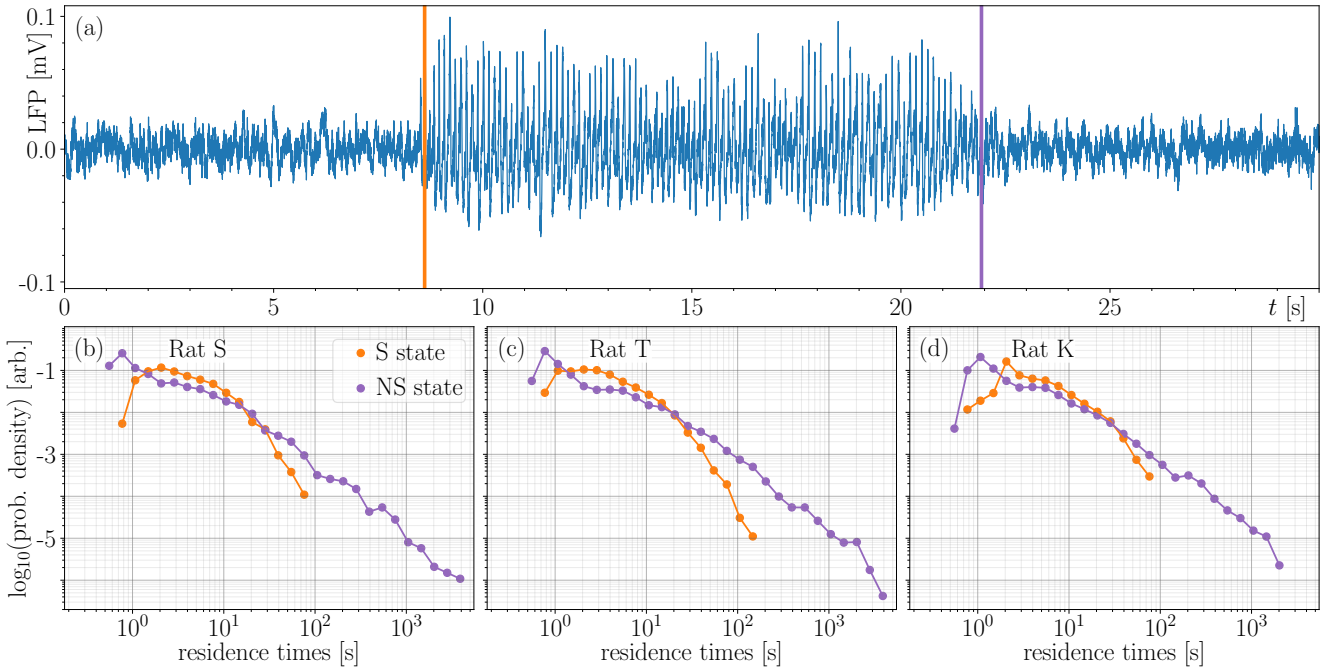

Figure 1. Characteristics of real seizure activity in GAERS. (a) An example of voltage recordings of real seizure activity, where the vertical lines indicate times of electrical seizure (orange) onset and (purple) offset according expert annotations. (b)-(d) The probability density of residence times in the S and NS states of rat S, T, and K respectively.

CTs between the NS and S states in the model's output and in the voltage recordings, and (ii) train a machine learning classifier using the model's output and then apply it to the voltage recordings.

Figure 1 (a) shows a typical example of real seizure activity in GAERS, taken from recording session 'S5A' (rat S, day 5 of recording, A^{th} recording session). This voltage recording was obtained for $t \geq 0$ with time resolution of 0.001 seconds via the procedure outlined in M1 of the Methods section. The vertical orange and purple lines in Fig. 1 (a) indicate the electrical seizure onset and offset times, respectively, according to the expert. For the ith seizure, we denote the onset time by $t_1^{(i)} > 0$ and the offset time by $t_2^{(i)} > t_1^{(i)}$. The same notation is used to denote times where our algorithm detects CTs between the NS and S states in our mathematical model. Note that the (orange) electrical seizure onset is characterised by the sudden change in the voltage recordings from a small-amplitude and weakly-correlated oscillation to a larger-amplitude and highly-correlated oscillation. The opposite occurs for the (purple) electrical seizure offset. These sudden changes and the timescale of the oscillations are two important characteristics that we tune our mathematical model to mimic as closely as possible.

The seizure in Fig. 1 (a) is just one of several seizures that have been annotated by the expert in this recording session. Looking at all seizures in the session, there is significant variation in the length of the different intervals of real seizure activity ($t_2^{(i)} - t_1^{(i)}$), which we refer to as *residence times in the S state,* and real non-seizure activity ($t_1^{(i+1)} - t_2^{(i)}$), which we refer to as *residence times in the NS state*. Figures 1 (b)-(d) shows the probability density of residence times in the S and NS states from 'rat S', 'rat T', and 'rat K', using all the annotations of electrical seizure onset and offset times provided by the expert. The information presented here is based on 621, 1593, and 1143 seizures from each respective rat. Note that the residence times in the S state are noticeably shorter than those in the NS state, and that residence times for both states have a similar range for different rats. The range of residence times is the third important characteristic that we tune our mathematical model to mimic for noise-induced CTs.

A canonical mathematical model

In most mathematical models of seizure dynamics, NS states are represented by noisy oscillations about a stable equilibrium point near 0 mV, and S states are represented by noisy motion along a non-stationary stable state, such as a stable limit cycle, with amplitude much larger than the amplitude of the noisy oscillations near 0 mV. If the model is well designed, NS states mimic the small-amplitude and weakly-correlated oscillations, and S states mimic the larger-amplitude and highly-correlated oscillations seen in voltage recordings such as Fig. 1. These mathematical models can be separated into the following groups: process-based models which are derived from a number of physical processes occurring in the brain to produce seizure-like

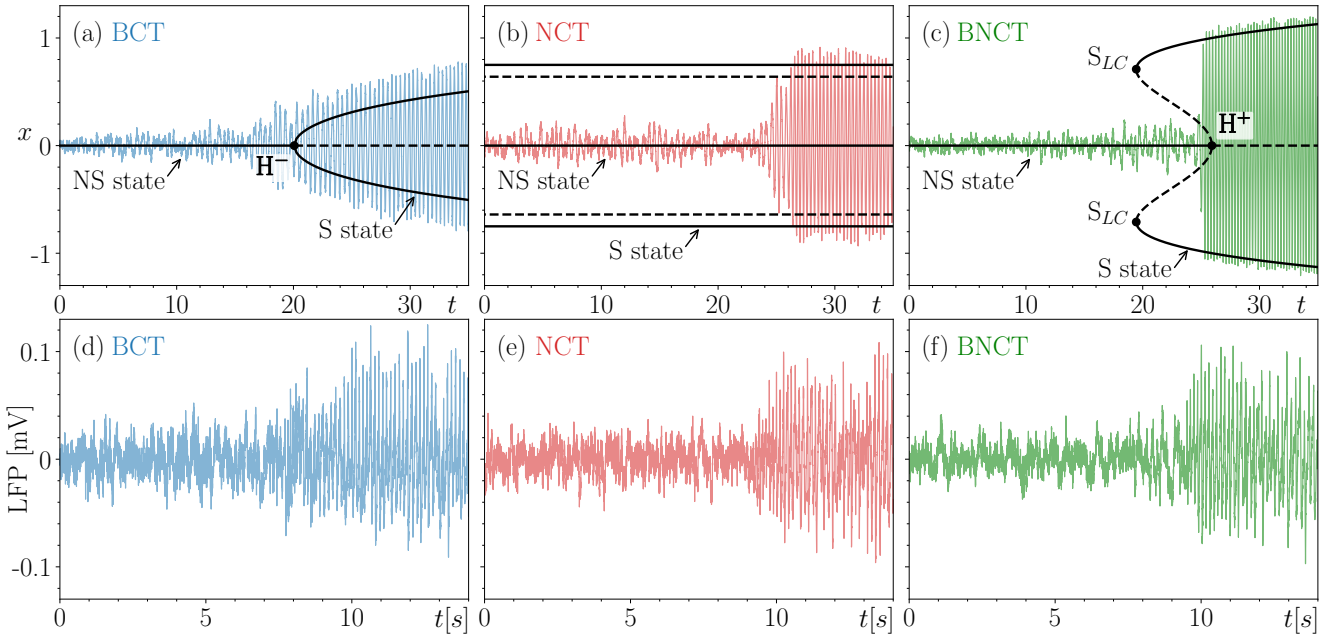

Figure 2. Examples of the three CT-types from the NS to S state in the model's output. (a) Bifurcation-induced CT due to drifting across a supercritical bifurcation point H^-. (b) Noise-induced CT in the bistability region between the stable NS and S states. (c) Bifurcation/noise-induced CT due to a combination of drifting across a subcritical bifurcation point, H^+, and bistability. The solid (dashed) lines represent the stable (unstable) NS and S states of the noise-free system (1), and S_{LC} in (c) denotes the fold of limit cycles bifurcation point. The coloured time series show the outputs of our model (5). Lower row shows corresponding CTs in voltage recordings classified via our framework (see Fig. S-20 in [30] for further details).

activity[16,23–25], phenomenological models which are designed to mimic real seizure activity in voltage recordings without taking into account the underlying processes in the brain[11–13,26–28], and a mixture of both[29]. CTs from the NS to S state in these models can be related to processes in the brain and largely categorised according to three CT-types, namely:

- Bifurcation-induced CTs (BCTs): The brain begins in the NS state, which is the only stable state. However, the NS state becomes unstable when the brain slowly drifts across a supercritical Hopf bifurcation point due to changing environmental or structural parameters. This causes the brain to transition to the stable S state, which gradually emerges from the bifurcation point and is the only stable state beyond it. See Fig. 2 (a).
- Noise-induced CTs (NCTs): The brain is bistable, meaning that it has two stable states, NS and S, that coexist. A sudden transition from the NS to S state is triggered by noise alone, in the form of endogenous or exogenous disturbances. See Fig. 2 (b).
- Bifurcation/noise-induced CTs (BNCTs): The brain is bistable as it slowly drifts towards a subcritical Hopf bifurcation, at which point the NS state loses stability and the S state becomes the only stable state. A combination of noise and the imminent instability of the NS state causes a sudden transition to the S state. See Fig. 2 (c).

We now introduce a phenomenological model in the form of two coupled stochastic differential equations. This model has equivalent NS and S states and mimics the three important characteristics of the voltage recordings specified in the previous section. Most importantly, it can also generate the three CT-types listed above. In the first step, we introduce and analyse the noise-free version of the model to identify candidates for NS and S states, as well as CTs between them. In the second step, we include random noise for direct comparison with voltage recordings of real seizure activity.

We consider the Bautin (generalised Hopf) bifurcation normal form[31] for the time evolution of a complex-valued variable z that we augment with a shear term,

$$\frac{dz}{dt} = \gamma \left((\mu + i\omega)\, z + s\,(1 + i\sigma)\, |z|^2 z - b\,|z|^4 z \right), \tag{1}$$

where $\mu \in \mathbb{R}$ and $s \in \mathbb{R}$ are bifurcation parameters, $\omega \in \mathbb{R}$ is the frequency of small-amplitude oscillation, $\sigma \in \mathbb{R}$ in the augmented term is the shear parameter, and $b = 1$ remains fixed throughout our investigations and is no longer explicitly referred

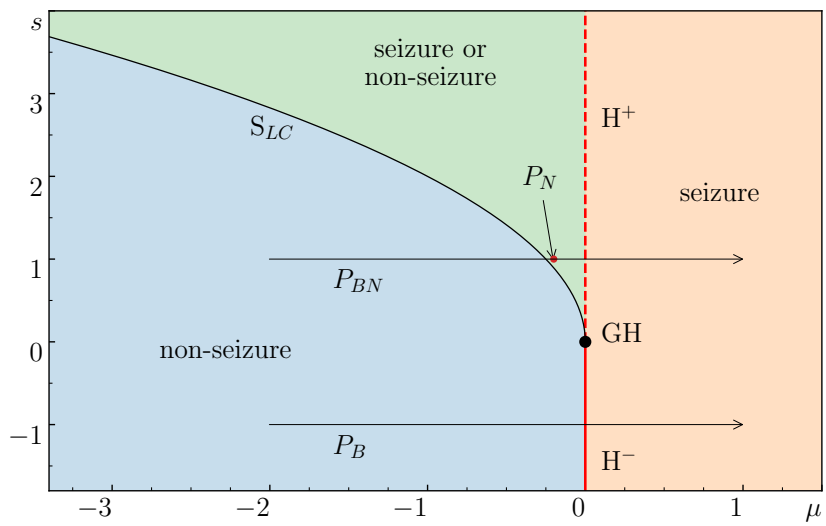

Figure 3. Two-parameter stability diagram for System (2) in the (μ, s) parameter plane showing the regions of (blue) stable NS state, (red) stable S state, and (green) bistability between the stable NS and S states. S_{LC} denotes the curve of saddle-node of limit cycle bifurcations, H^- and H^+ denote the curves of supercritical and subcritical Hopf bifurcations, respectively, and GH denotes the Bautin (generalised Hopf) bifurcation point. The parameter paths, P_B and P_{BN}, and the point P_N, denote the values of μ and s used to generate BCTs, BNCTs, and NCTs, respectively, in the model's output. Note that the bifurcation curves do not depend on any other system parameters.

to. The timescale parameter $\gamma > 0$ allows us to adjust the timescale of the solutions. For a fixed s we choose σ so that these parameters are always of the same sign. Although the shear term is not present in normal forms, it is essential for analysing a system's response to external disturbances[32–34]. In our mathematical model, it enables us to produce seizure-like activity that more closely resembles real seizure activity seen in the voltage recordings.

We now discuss the existence of candidates for NS and S states and then conduct linear stability and bifurcation analysis to uncover different parameter paths in the (μ, s) parameter plane that give rise to the different types of CT between these states. To simplify the discussion, we transform Eq. (1) to a polar coordinate system by setting $z = re^{i\theta}$ where $r \geq 0$ and $\theta \in [0, 2\pi)$,

$$\begin{aligned} \frac{dr}{dt} &= \gamma r \left(\mu + s r^2 - r^4\right), \\ \frac{d\theta}{dt} &= \gamma \left(\omega + \sigma s r^2\right). \end{aligned} \tag{2}$$

Note that the dr/dt equation is decoupled from the $d\theta/dt$ equation in the sense that it does not depend on θ and can thus be solved on its own. An equilibrium point r_e for the dr/dt equation corresponds to an equilibrium point in the full system if $r_e = 0$, and to a limit cycle in the full system if $r_e > 0$. Thus, System (2) has an equilibrium point, e_1, that is located at the origin and exists for all parameter values. Noisy motion around e_1 corresponds to an NS state. Furthermore, System (2) has two limit cycles, L_+ and L_-, whose radii are given by

$$r_+ = \sqrt{\frac{1}{2}\left(s + \sqrt{s^2 + 4\mu}\right)} \quad \text{and} \quad r_- = \sqrt{\frac{1}{2}\left(s - \sqrt{s^2 + 4\mu}\right)}, \tag{3}$$

respectively. L_+ exists for $\mu > -s^2/4$ if $s > 0$ and for $\mu > 0$ if $s < 0$, whereas L_- exists for $-s^2/4 < \mu < 0$ and $s > 0$. Noisy motion around L_+ corresponds to an S state.

Linear stability analysis of equilibria for the dr/dt equation reveals the following. The equilibrium e_1 is stable for $\mu < 0$ and unstable for $\mu > 0$. The limit cycle L_+ is stable while L_- is unstable. Bifurcations of e_1, L_-, and L_+ for changes in μ and s are shown in a two-parameter bifurcation diagram in Fig. 3. For $s < 0$ and as μ changes from negative to positive values, e_1 becomes unstable and L_+ is born in a supercritical Hopf bifurcation H^- that occurs when crossing the solid part of the vertical red line at $\mu = 0$; see also the black solution branches in Fig. 2 (a). For $s > 0$ and as μ changes from positive to negative values, e_1 becomes

	BCT	BNCT	NCT
μ: bifurcation parameter	$\mu(t) = -2 + t/20$ for $t \in [0, 60]$	$\mu(t) = -2 + t/20$ for $t \in [0, 60]$	-0.22
s: bifurcation parameter	-1	1	1
σ: shear	-1	1	1
ν: noise level	0.18	0.18	0.18
ω: oscillation freq.	1.3	1.3	1.3
γ: time scale	10	10	10

Table 2. The parameter values chosen for our model in Eqs. (5) to generate different CT-types.

stable and the unstable L_- is born in a subcritical Hopf bifurcation H^+ that occurs when crossing the dashed part of the vertical red line at $\mu = 0$. By decreasing μ further, L_+ and L_- collide and disappear in a saddle-node of limit cycles bifurcation S_{LC} that occurs when crossing the black solid curve given by $\mu = -s^2/4$ and $s > 0$; see also the black solution branches in Fig. 2 (c). Note that e_1, L_+, and L_- coincide at a special Bautin (generalised Hopf) bifurcation point GH at $(\mu, s) = (0, 0)$.

As emphasised by the different coloured regions in Fig. 3, the bifurcations of e_1, L_-, and L_+ define three qualitatively different types of activity. In the blue region, where e_1 is the only stable state, there is non-seizure-like activity. In the orange region, where L_+ is the only stable state, there is seizure-like activity. In the green region of bistability, where both e_1 and L_+ are stable, either non-seizure-like or seizure-like activity can occur.

In Fig. 3 we also indicate two parameter paths and one special point. The paths P_B and P_{BN} are used to generate many BCTs and BNCTs from the NS to S state, respectively. This is done by repeatedly increasing μ from -2 to 1 for $s = 1$ and $s = -1$, respectively, according to

$$\mu(t) = -2 + t/20, \tag{4}$$

for $t \geq 0$. The point P_N at $s = 1$ and $\mu = -0.22$ is used to generate a long time series containing many NCTs between the NS and S states.

In order to provide a direct comparison with the real seizure activity seen in voltage recordings such as Fig. 1, we transform Eq. (1) into its Cartesian form by setting $z = x + i\,y$, where x and y are real-valued variables, and further augment it by including additive noise terms to generate the three different CTs mentioned above,

$$\begin{aligned} \frac{dx}{dt} &= \gamma \left(\mu x + s \left(x^2 + y^2\right) (x - \sigma y) - \omega y - \left(x^2 + y^2\right)^2 x \right) + \nu \eta_x(t), \\ \frac{dy}{dt} &= \gamma \left(\mu y + s \left(x^2 + y^2\right) (y + \sigma x) + \omega x - \left(x^2 + y^2\right)^2 y \right) + \nu \eta_y(t). \end{aligned} \tag{5}$$

Here, $\eta_x(t)$ and $\eta_y(t)$ are independent Gaussian random variables (white Gaussian noise with zero mean and unit variance) and ν is the noise level (standard deviation of the noise). To generate solutions for x and y at a given time t, we discretise System (5) according to the Euler-Maruyama method for a fixed time step δ_t. In our numerical experiments, we set $\delta_t = 0.001$, $x(0) = x_0 = 0.1$, and $y(0) = y_0 = 0.1$. This value of δ_t is chosen to prevent numerical errors from arising when σ is relatively large. The choice of initial condition (x_0, y_0) ensures that all simulations of System (5) begin from the NS state.

Throughout the rest of the paper, we refer to System (5) as the 'model', and to a time series of $x(t)$ as the 'model's output'. Furthermore, the units of quantities related to voltage recordings are millivolts (mV) and time is in seconds (s), whereas the units of quantities related to the mathematical model are arbitrary. This applies to quantities specified in the text, in tables, and in figures and their captions.

Seizures in the model's output resemble voltage recordings of real seizure activity

Earlier in the paper, we identified three characteristics of the real seizure activity that we want our mathematical model to mimic:

(i) The amplitude and correlation of oscillation in the NS and S states.

(ii) The range of residence times in the NS and S states.

(iii) Intrinsic time scales of the NS and S states.

When tuning the model parameters we found that trade-offs emerge and we have to make choices on what characteristics should be prioritised. Optimising μ for characteristic (i) makes it difficult to find a ν to mimic characteristic (ii). Similarly, optimising γ for characteristic (iii) makes it difficult to mimic characteristic (ii), and vice-versa. We find that characteristics (i) and (ii) are the most critical for training the SVM classifier to accurately classify seizure generation mechanisms in time series. We therefore choose to prioritise characteristics (i) and (ii) and accept a trade-off in characteristic (iii). The model parameters in Table 2 have been chosen accordingly.

Characteristic (i): The top row in Fig. 4 shows that the model's output (a) closely resembles real voltage recordings in GAERS (b). Characteristic (i) is reasonably well satisfied since the relative change in the amplitude of oscillation for a CT from the NS to S state in the model's output is similar to the seizure onset in the voltage recordings. The same can be said forf the CT from the S to NS state and the seizure offset. Figures 4 (a) and (b) also highlight one benefit of including shear, quantified by σ in the model: the model exhibits variations in local maxima/minima observed in real voltage recordings beyond what can be achieved with noise alone.

Characteristic (ii): In the bottom row, Fig. 4 (e) shows the probability density of residence times in the S state for rat S (blue) and the model (orange). Figure 4 (f) shows the same in the NS state. These residence times were obtained by applying our CT detection algorithm, introduced in the next subsection, to the voltage recordings for rat S, and to the model's output containing several NCTs between the NS and S states. Figures 4 (e) and (f) show that the model exhibits a range of residence times in the NS and S state, which is the same or wider than those in the voltage recordings. However, the model tends to spend longer in the S state than the voltage recordings, resulting in some differences in the magnitude of the probability density.

Characteristic (iii): In the middle row, Figs. 4 (c) and (d) show spectrograms of the time series from Figs. 4 (a) and (b), respectively. While in the S state, the model's output has a component at the same frequency as the voltage recording. However, the dominant component in the voltage recording is about twice that of the model's output.

Step 2: Detecting CTs and identifying distinctive properties of each CT-type in the model's output

The aim of this paper is to classify the types of CT that are responsible for seizure onset in voltage recordings of real seizure activity using a classifier trained on the model's output. To ensure consistency, we now develop a single algorithm that detects CTs between the NS and S states in both the voltage recordings and the model's output.

To detect CTs in a noisy time series, one typically specifies a threshold and defines a CT to a different state as a crossing of the threshold. However, such a single threshold approach has its shortcomings. First, a single threshold can be crossed multiple times in a short time interval due to noise, leading to false CTs. Second, if one of the states is oscillatory, the threshold can be crossed multiple times while the system is in that state, also leading to false CTs. Both shortcomings arise in the case of CTs to the S state and we address them as follows:

- We use the concept of a 'non-ideal relay' with two thresholds[35] and define CTs between the NS and S states in terms of successive crossings of the two thresholds.
- We use a 'moving window analysis' to assess how long the system remains in the new state after successive crossings of the two thresholds.

Constructing a CT detection algorithm

The time resolution of the voltage recordings determines the time step size δ, i.e., the time interval between two consecutive points in the time series. The model's output is generated with the same time step size. In the case of voltage recordings in GAERS, we have $\delta = 0.001$.

Additionally, we introduce six algorithm parameters: the "on" threshold $\alpha > 0$, the "off" threshold $0 < \beta < \alpha$, the size of the moving window $\tau_w > 0$, the time step size of the moving window $\delta \leq \Delta \leq \tau_w$, the minimum time duration of larger-amplitude oscillations, $\tau_S > \tau_w$, expressed as $\tau_S = n_S\Delta + \tau_w$, where n_S is an integer, and the minimum time duration of lower-amplitude oscillations, $\tau_{NS} \geq \tau_S$, expressed similarly as $\tau_{NS} = n_{NS}\Delta + \tau_w$, where n_{NS} is an integer. We list these parameters and specify their values in Table 3.

The starting point: We start in the NS state, where $|x(t)| < \beta$ for the time duration of at least τ_{NS}.

The moving window: When the system is in the NS state and $|x(t)|$ exceeds α at time $t = t_j$, the moving window is activated and $|x(t)|$ is examined within consecutive windows of duration τ_w that are shifted in time by Δ, starting with $[t_j, t_j + \tau_w]$, then $[t_j + \Delta, t_j + \tau_w + \Delta]$, $[t_j + 2\Delta, t_j + \tau_w + 2\Delta]$, and so on. The moving window is deactivated in two cases: (i) the system is in the NS state, the window is activated, but the algorithm does not detect a CT to the S state, and (ii) the system is in the S state

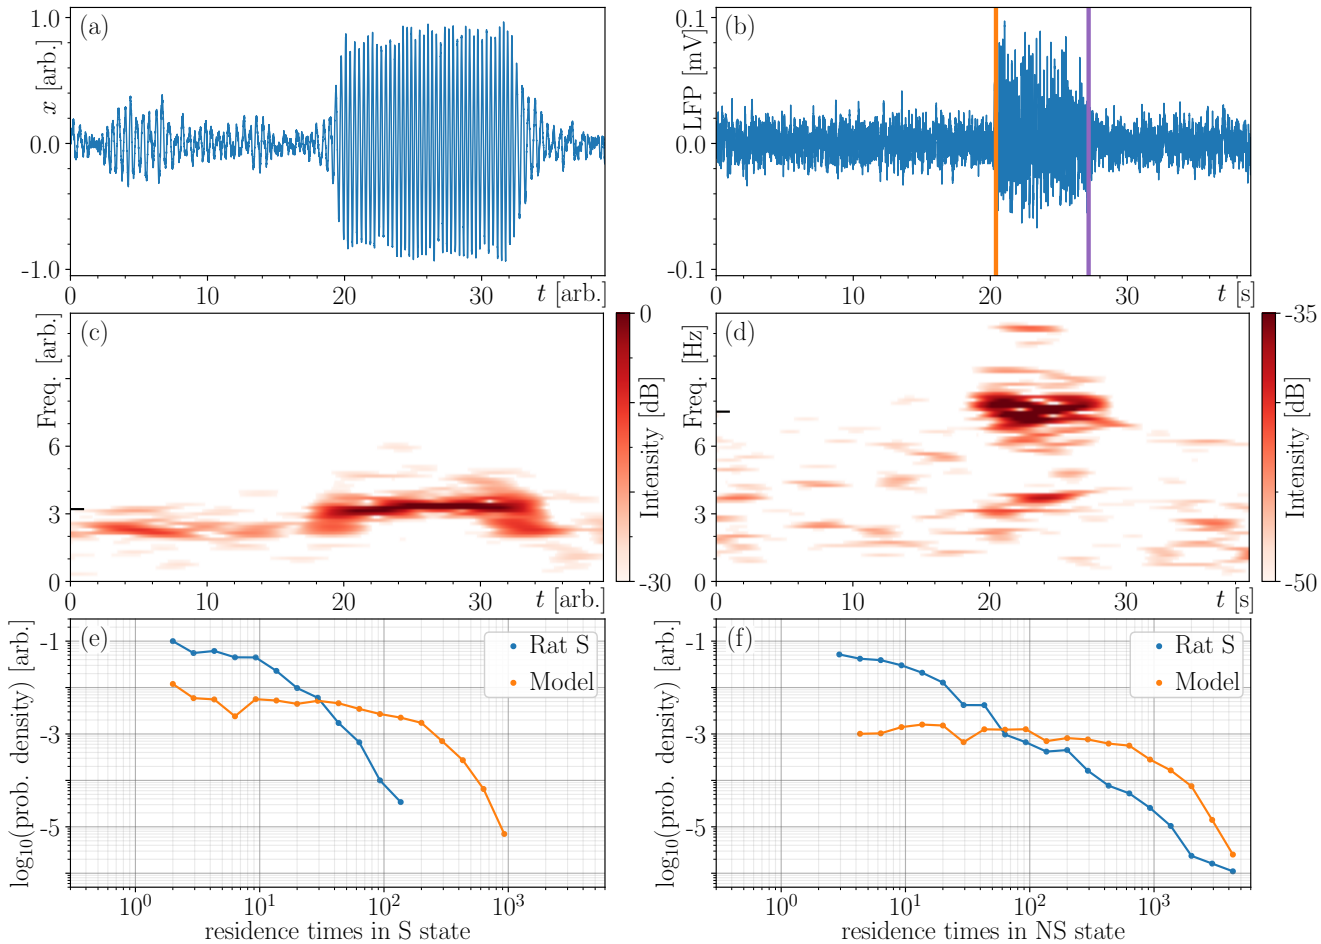

Figure 4. A comparison of the characteristics of (a) the model's output and (b) the actual voltage recordings for rat S. (c)-(d) The spectrograms of the time series from (a) and (b), respectively. (e)-(f) The probability density of residence times in the S and NS state, respectively, for (blue) rat S and (orange) the model's output.

and the algorithm detects a CT to the NS state.

Critical transitions: The algorithm detects a *CT from the NS to S state* at time $t = t_1$ if:

(a1) The system is in the NS state just before t_1.

(a2) $|x(t)|$ exceeds α at time $t = t_1$, i.e., $|x(t_1)| = \alpha$ and $|x(t_1 + \delta)| > \alpha$.

(a3) Each of the n_S consecutive positions of the moving window contains an $|x(t)| > \beta$.

The algorithm detects a *CT from the S to NS state* at time $t = t_2$ if:

(b1) The system is in the S state just before t_2.

(b2) $|x(t)|$ falls below β at time $t = t_2$, i.e., $|x(t_2)| = \beta$ and $|x(t_2 + \delta)| < \beta$.

(b3) Each of the n_{NS} consecutive positions of the moving window contains no $|x(t)| \geq \alpha$.

In other words, the algorithm detects a CT from the NS to S state if $|x(t)|$ exceeds the upper threshold α and then continues to exceed the lower threshold β frequently enough for a period of at least τ_S. Similarly, the algorithm detects a CT from the S to NS state if $|x(t)|$ falls below the lower threshold β and then does not exceed the upper threshold α for a period of at least τ_{NS}. In Fig. 5 we provide an illustrative example of how the above algorithm detects CTs in the model's output from Fig. 4 (a) using the algorithm parameters specified in Table 3.

Almost-occurring critical transitions: The algorithm detects an *almost-occurring CT from the NS to S state* at time $t = \tilde{t}_1$ if

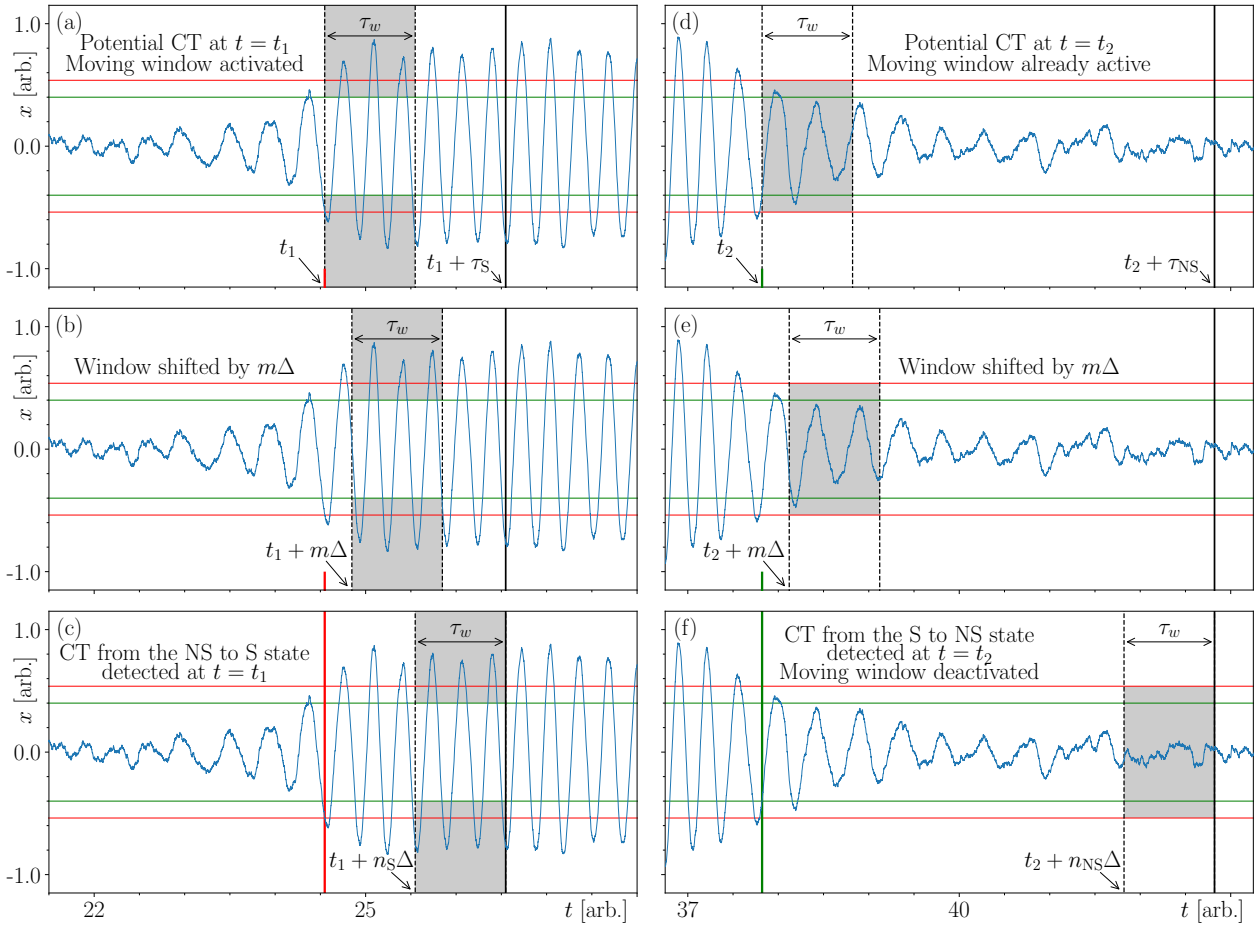

Figure 5. Illustrating how the CT detection algorithm works using the model's output from Fig. 4 (a) as an example. (a)-(c) Detecting a CT from the NS to S state at time $t = t_1$. (d)-(f) Detecting a CT from the S to NS state at time $t = t_2$. The detection algorithm parameters are specified in Table 3, and we used $m = 300$ in (b) and (e). The red and green horizontal lines indicate the thresholds α and β, respectively.

(a1) and (a2) are satisfied but (a3) is not. Similarly, the algorithm detect an *almost-occurring CT from the S to NS state* at time $t = \tilde{t}_2$ if (b1) and (b2) are satisfied but (b3) is not.

Applying the CT detection algorithm to the model's output and voltage recordings

We refer to S1 and S2 in [30] for examples of CTs and almost-occurring CTs detected by the algorithm in both the model's output and the voltage recordings, and examples of artefacts in the voltage recordings and how we prevent the algorithm from detecting these as CTs. In S3 and S4 in [30] we show how the shear parameter σ influences the detected residence times in the NS and S states and the times at which the algorithm detects CTs in the model's output. In S5 in [30] we describe how the algorithm parameters are tuned to maximise the agreement between the CTs detected by the algorithm and the CTs annotated by the expert in the voltage recordings. These parameters are specified in Table 3.

Identifying distinctive properties of the model's output that differentiate between the three CT-types

When a system is perturbed away from its stable state — whether stationary or not — it will attempt to return to this state. However, the closer the system is to a bifurcation point at which the state loses stability, the less stable it becomes and the longer it takes to return. Therefore, as a noisy system gradually approaches a bifurcation point, one would expect to see an increase in the variance and autocorrelation of the observed time series. This phenomenon is known as 'critical slowing down' (CSD) and the accompanying increases in the variance and autocorrelation are commonly referred to as 'early warning signals' (EWSs). Scheffer *et al.*[9] recognised that EWSs can be detected in time series recordings of a single observable of a complex system, allowing researchers to detect CSD without knowledge of the entire state of the system. This approach has been widely successful in studies of climate and environmental systems, enabling researchers to forecast when BCTs may occur[36,37]. For

Time series and detection algorithm parameters	Model's output	Voltage recordings
δ: time step size	0.001	0.001
α: on threshold	0.55	[0.03, 0.1]
β: off threshold	0.45	α - 0.01
τ_w: size of the moving window	1	1
Δ: time step size of the moving window	0.001	0.001
τ_{S}: minimum time duration of the S state	2	2
τ_{NS}: minimum time duration of the NS state	5	3

Table 3. The values of the time step size and CT-detection algorithm parameters used for the model's output and voltage recordings. α is chosen differently for each voltage recording session according to the method outlined in S5 in [30].

recent publications on this area we refer to Dakos *et al.*[17] and Ashwin *et al.*[38]. Furthermore, in terms of classifying CTs in a noisy time series, Bury *et al.*[39] classified different bifurcation-induced CTs based on characteristic increases in the variance and autocorrelation of the time series prior to the different CTs.

CTs in the brain occur on timescales that are several orders of magnitude faster than those in climate and environmental sciences and could also be noise-induced. Therefore, our approach differs from previous work on EWSs in three key ways: First, we analyse the time series around a CT, that is, shortly before and shortly after a CT, rather than just prior to a CT. This is because our primary goal is to identify the CT-type rather than to foresee one that is about to occur within seconds. (We will discuss EWSs for the onset of seizures and their predictive power in future work.) Second, we consider the following properties of the time series and their slopes around a CT from the NS to S state; see also Table 5:

(TSP1) Gaussian variance (GV).

(TSP2) Base-10 logarithm of the Gaussian variance ($\log_{10}$GV).

(TSP3) Lag-1 autocorrelation (AC).

(TSP4) Base-10 logarithm of the Gaussian variance of the lag-1 autocorrelation ($\log_{10}$GV(AC)).

We collectively refer to (TSP1)-(TSP4) as the 'time series properties' and denote them by TSPs. We also consider the 'slopes of the TSPs', and denote them by m(TSPs, t_m), where t_m is the chosen time interval over which the slope is calculated. In other words, the m(TSPs, t_m) carry information about the trends in TSPs over the past t_m seconds. Third, and most importantly, we use the TSPs and m(TSPs, t_m) to classify CTs from the NS to S state as (i) bifurcation-induced, (ii) noise-induced, or (iii) bifurcation/noise-induced.

Selection criteria for CTs: We select 100 examples of each CT-type from the NS to S state in the model's output according to the following criteria, bearing in mind that the algorithm detects each CT at a time $t = t_1$,

- The model is in the NS state for all $t \in \left[t_1 + T^-, t_1\right)$, where $T^- = -30$.
  $|T^-|$ is the length of the time series considered prior to each detected CT.
- The model is in the S state for all $t \in \left[t_1, t_1 + T^+\right]$, where $T^+ = 10$.
  T^+ is the length of the time series considered after each detected CT.

Note that these criteria exclude some CTs detected by the algorithm. For instance, if insufficient time is spent in the NS state prior to a CT, the CT is excluded; see S10 in [30] for further information. Furthermore, while our CT-type classifier requires sufficient time prior to a CT, there is also a biological reason to exclude such CTs: the brain may not have sufficient time to fully return to the NS state beyond what can be seen in the voltage recording.

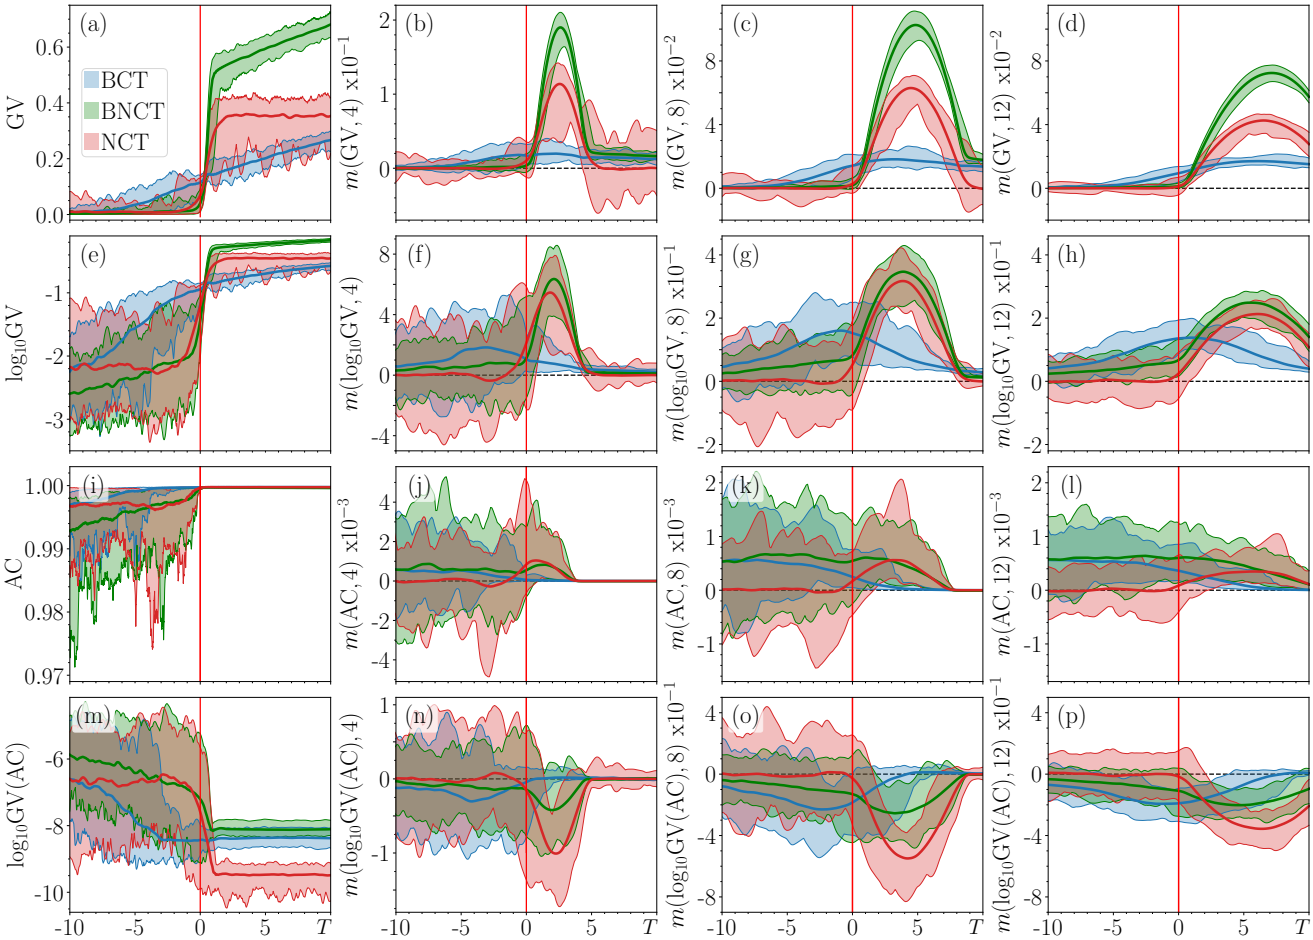

Figure 6. Mean values (solid curves) and range of values (shaded regions surrounding each curve) of the different TSPs and m(TSPs, t_m) (specified on each vertical axis) before and after (in blue) BCTs, (in red) NCTs, and (in green) BNCTs plotted versus $T = t - t_1$, for the t_1 detected by our algorithm for each CT. The first column corresponds to the TSPs, the remaining three columns correspond to the m(TSPs, t_m) for $t_m = 4$, 8, and 12 respectively.

Computing TSPs and m(TSPs, t_m) near selected CTs: When computing TSPs and m(TSPs, t_m) near each CT, it is convenient to work with a new time, $T = t - t_1$, where $T \in \left[T^-, T^+\right]$ and each CT is detected at $T = 0$. A single value of each TSP is computed over a time interval denoted by t_w and referred to as the 'window length'. We consider a fixed $t_w = 1$ for both the model's output and voltage recordings; see Fig. S-12 in [30] for justification. A single value of each m(TSPs, t_m) is computed over a longer time interval, $t_m > t_w$, referred to as the 'slope length', and we examine different values of t_m. Thus, for a given choice of T^- and T^+, we obtain values of (TSP1)-(TSP3) in the interval $\left[T^- + t_w, T^+\right] = \left[T^- + 1, T^+\right]$, (TSP4) in the interval $\left[T^- + 2t_w, T^+\right] = \left[T^- + 2, T^+\right]$, and all m(TSPs, t_m) in the interval $\left[T^- + 2t_w + t_m, T^+\right] = \left[T^- + 2 + t_m, T^+\right]$; see M2 of the Methods section and S6 in [30] for additional details.

We obtain all the TSPs and m(TSPs, t_m) for the 100 selected examples of each CT-type, summarising the results in Fig. 6. We plot the mean values (solid curves) and spread (shaded bands) as a function of $T \in [-10, 10]$, where the most prominent changes occur, and use a different colour for each CT-type. More specifically, in the first column of Fig. 6 we plot the mean value taken by each of the TSPs versus T. In the remaining three columns, we plot the mean value taken by each of the m(TSPs, t_m) versus T for $t_m = 4$ (second column), $t_m = 8$ (third column), and $t_m = 12$ (fourth column). In Fig. S-11 in [30], we show how the TSPs behave around a single example of each CT-type.

Differentiating between the three CT types based on TSPs and m(TSPs, t_m): While some aspects of Fig. 6 are similar to previous findings by Milanowski and Suffczynski[12], we focus on how the TSPs and m(TSPs, t_m) can be used to distinguish between the three CT-types rather than identify any EWSs, which we leave for future work. For instance, we see in Fig. 6 (c) that three separate (non-overlapping) bands of m(GV, t_m) emerge near $T = 5$, where each band corresponds to a different CT-type. This allows one to visually distinguish between each CT type. Comparing Figs. 6 (b), (c), and (d), we see that as t_m increases,

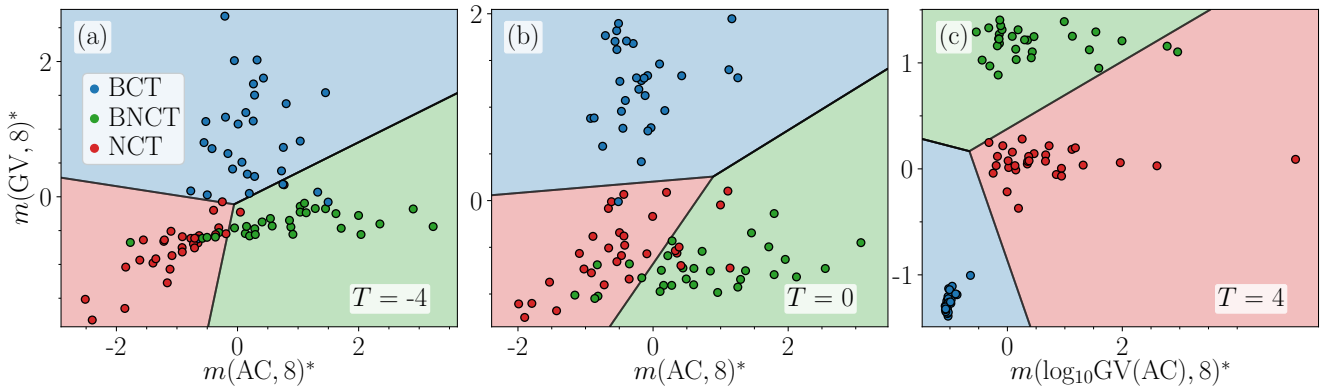

Figure 7. An illustrative example of how a two-feature SVM works, with the features specified on each axis. The different coloured points correspond to (blue) the true BCTs, (green) the true BNCTs, and (red) the true NCTs. This information is known to us but not to the SVM. The three different coloured regions are where the SVM classifies CTs as (blue) BCTs, (green) BNCTs, and (red) NCTs. The accuracy of the classification depends on the chosen features and whether classification is performed (a) before a detected CT, (b) at a detected CT, or (c) after a detected CT.

the width of these bands decreases, thereby improving our ability to visually distinguish between different CT types. As another example, we can distinguish between the three CT types by combining the following two observations when $T > 2$: BNCTs can be distinguished from BCTs and NCTs based on the separation between corresponding GV bands in Fig. 6 (a), and NCTs can be distinguished from BCTs and BNCTs based on the separation between corresponding $\log_{10}$GV(AC) bands in Fig. 6 (m). By considering several TSPs and m(TSPs, t_m) simultaneously we may distinguish between the three CT-types at earlier T values, even before a CT.

Step 3: Training and optimising a Support Vector Machine on CTs of known type in the model's output. It is convenient to automate the process of differentiating between the three CT-types based on their distinctive TSPs and m(TSPs, t_m) using a machine learning classifier. We chose a Support Vector Machine (SVM) because it is an explainable machine learning technique, i.e., there is no ambiguity about how it works. An SVM provides an optimal separatrix (e.g. a line or a hyperplane) that separates each class (in our case, the CT-type) in an F-dimensional space, where $F \in \mathbb{N}^+$ is the number of 'features'. In other words, the F-dimensional space is divided into different regions that correspond to each class and the feature data is classified according to which region of the F-dimensional space it resides in.

Illustrating how SVMs work: For illustrative purposes, we consider a simple SVM with only two features ($F = 2$), to show how CTs in the model's output are classified and how the classification changes for different values of T. The three panels in Fig. 7 are obtained for $T = -4$ in (a), $T = 0$ in (b), and $T = 4$ in (c). The features in each panel are chosen from Fig. S-13 (b) in [30] as the two most important features at these values of T. Each panel is divided into three regions separated by the optimal lines, and these regions are coloured according to each classification class: blue for BCT, green for BNCT and red for NCT. The different coloured points correspond to the true BCTs in blue, true BNCTs in green, and true NCTs in red. These are known to us but not to the SVM. The $*$ superscript on each axis label indicates that the corresponding quantities have been scaled as described in M3 of the Methods section.

Figure 7 shows that even a two-feature SVM can distinguish between seizure generation mechanisms in noisy time series with less CTs misclassified as T increases. Figure 7 (a) shows that shortly before a CT, when $T = -4$, a small number of BCT points are misclassified as BNCTs, some BNCT points are misclassified as NCTs, and only two NCT points are misclassified. Figure 7 (b) shows that at a CT, when $T = 0$, there is a smaller number of misclassifications of NCT and BNCT points, while only one BCT point is misclassified. Figure 7 (c) shows that shortly after a CT, when $T = 4$, only two BNCT points are misclassified as NCTs.

Training SVMs: We use the TSPs and m(TSPs, t_m) of the model's output, where we know how each CT is generated, as features to train different SVMs to classify the three CT-types from the NS to S state. To determine how classifications vary depending on what combination of features are used, we train three types of SVM:

- **Type-1** using the four TSPs, so that $F = 4$.
- **Type-2** using the four m(TSPs, t_m), so that that $F = 4$.
- **Type-3** using the four TSPs and the four m(TSPs, t_m), so that $F = 8$.

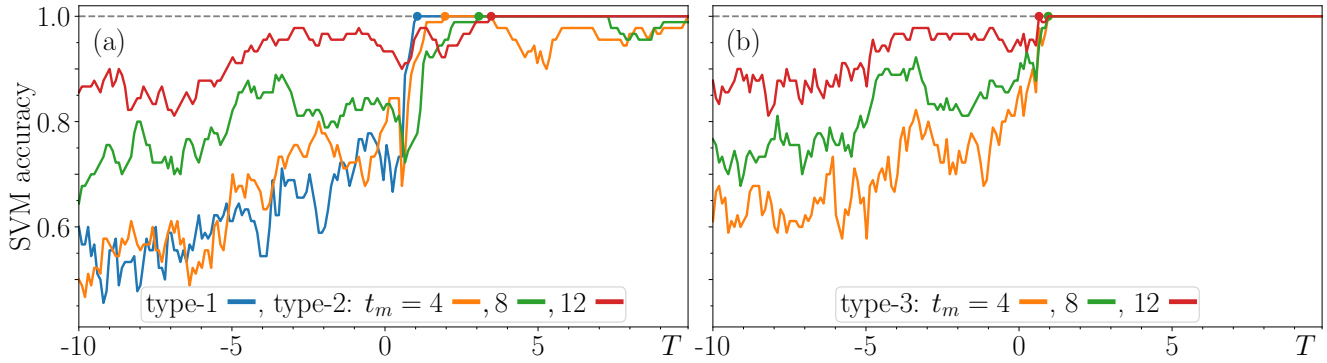

Figure 8. SVM accuracy in classifying unseen data from the model's output vs. T for (a) SVM type-1 and 2 and (b) SVM type-3. For SVM type-2 and 3, we include the dependence on different values of t_m specified in the plot legends. The dots indicate when each SVM achieves perfect accuracy for the first time.

The training procedure is the same for each SVM and is outlined in M3 of the Methods section. The combinations of features used for each classifier type are listed in Table 6 as a point of reference.

Testing and optimising SVMs: We test each SVM by examining its ability to classify CTs in unseen data from the model's output. This helps us decide which SVM type is most suitable for classifying seizure generation mechanisms in voltage recordings. SVM performance is quantified as the fraction of correctly classified samples. We refer to this quantity as the *SVM accuracy*, and examine how it changes with T for each SVM type. When the SVM accuracy reaches 1, we say that the SVM achieves perfect accuracy at this value of T. In Fig. 8 we plot the SVM accuracy versus T for SVMs type-1 and 2 in (a) and SVM type-3 in (b). The value of T at which a given SVM achieves perfect accuracy for the first time is indicated by a coloured dot. An additional analysis of 'feature importance' in each SVM type is provided in Figs. S-13 and S-14 in [30].

Figure 8 (a) shows that SVM type-1 performs poorly for most $T < 1$, achieves perfect accuracy before SVM type-2 for each t_m, and maintains its perfect accuracy for $T > 1$. Figures 8 (a) and (b) show that increasing t_m increases the accuracy of SVMs type-2 and 3, however, SVM type-2 requires larger T values to achieve perfect accuracy for the first time for larger t_m. Additionally, while SVM type-3 maintains its perfect accuracy after first achieving it, SVM type-2 does not beyond a certain T and starts to decrease in accuracy much sooner (at values of T closer to 0) when smaller t_m is used. Figures 8 (a) and (b) also show that SVM type-3 consistently outperforms SVM type-2 for the same values of T and t_m. It is clear that SVM type-3 benefits from the larger accuracy of SVM type-2 for $T < 1$ and the consistently perfect accuracy of SVM type-1 for $T > 1$. Crucially, Fig. 8 shows that we do not need to be as restrictive with our choice of T^- and T^+ as each SVM achieves perfect accuracy near $T = 0$. This result works to our advantage in the following subsection because choosing values of T^- and T^+ closer to 0 enables us to classify a larger number of CTs in the voltage recordings.

Based on the above discussion we conclude that SVM type-3 is the most suitable choice for classifying seizure generation mechanisms in the voltage recordings. However, an optimal choice of t_m must be made. Although SVM type-3 is more accurate for $T < 1$ when $t_m = 12$ than for $t_m = 8$, we can be less restrictive with our choice of T^- and T^+ for smaller t_m, allowing more CTs to be classified. This results in a trade-off between maximising either the accuracy or the number of CTs that can be classified. In the interest of classifying as many seizure generation mechanisms as possible with a reasonably high level of accuracy, we choose to use the SVM type-3 with $t_m = 8$ and refer to it as the 'model-trained SVM' to emphasise that it is trained on the model's output.

Step 4: Using model-trained CT-type classifier to classify unknown seizure generation mechanisms in the voltage recordings of real seizure activity

This subsection presents our main results: the classification of seizure generation mechanisms in voltage recordings of real seizure activity using the model-trained SVM. We use the same time notation as for the model's output to describe the voltage recordings near CTs from the NS to S state. Specifically, we use a new time $T \in [T^-, T^+]$, where $T = t - t_1$, and $T = 0$ is the time when the algorithm detects a CT from the NS to S state.

Stage 1 (Detection of CTs): We begin by using our detection algorithm to detect CTs between the NS and S states in the voltage recordings of rats S, T, and K. The algorithm parameters, specified in Table 3, are chosen according to the procedure outlined in S5 in [30]. We denote the total number of detected CTs from the NS to S state for a given rat by N_{det}, where 'det' stands for detected. We refer to this set of CTs as the 'detected CTs'.

Stage 2 (Filtering of detected CTs): We classify a detected CT only if it satisfies the following conditions:

C1: The brain remains in the NS state for all $T \in [T^-, 0)$.

C2: The brain remains in the S state for all $T \in \left[0, T^+\right]$.

C3: There is some overlap between what the algorithm identifies as the S state and the expert annotation of the S state.

C4: No artefacts appear for $T \in [T^-, 0)$.

C5: The threshold α is not crossed for any $T \in [T^-, 0)$.

The conditions C1 and C2 are consistent with those used when selecting CTs from the model's output and ensure that the analysed TSPs and m(TSPs, t_m) correspond to an individual CT. C3 ensures that we only classify CTs where our algorithm agrees with the expert annotation that there is a CT. The algorithm parameters used in Stage 1 are chosen to maximise this agreement as per the steps outlined in S5 in [30]. C4 and C5 exclude scenarios in which artefacts or almost-occurring CTs are found just prior to a detected CT, that is for $T \in [T^-, 0]$; see Figs. S-16 and S-21 in [30] for more details. The subset of detected CTs that satisfy all five conditions are referred to as the 'filtered CTs' and the number of filtered CTs for a given rat and choice of T^- and T^+ is denoted by $N_{\text{filt}}(T^-, T^+)$, where 'filt' is short for filtered.

Stage 3 (Classification of filtered CTs): For our choice of T^-, T^+, $t_w = 1$ and $t_m = 8$, we obtain all TSPs and m(TSPs, t_m) of the filtered CTs for $T \in \left[T^- + 2t_w + t_m, T^+\right] = \left[T^- + 10, T^+\right]$. This is the time interval where we can apply our model-trained SVM classifier to the filtered CTs. Next, we use $N_{\text{type}}(T)$ to denote the number of CTs classified at time T as type = BCT, BNCT or NCT , and compute the proportions of:

- Filtered CTs to the total number of detected CTs: $N_{\text{filt}}(T^-, T^+)/N_{\text{det}} \in [0, 1]$.
- Filtered CTs classified as a particular type at time T to the total number of filtered CTs: $N_{\text{type}}(T)/N_{\text{filt}}(T^-, T^+) \in [0, 1]$, bearing in mind that $T \in \left[T^- + 10, T^+\right]$.

We refer to Figs. S-17, S-18 and S-19 in [30] for details of classification dependence on T.

The main experiment

We now have all the ingredients needed to design our main experiment and address our main research question: What CT-types, or generation mechanisms, are responsible for seizure onset in GAERS. To demonstrate the robustness of our findings with respect to the selection of time intervals around the CTs, we consider a fixed value for T^+ and different values for T^-. More specifically, we set $T^+ = \tau_S = 2$ and consider $T^- = -16, -14, -12, -10$, and -8, which are chosen differently from the T^+ and T^- used for training the SVM classifier in order to to maximise $N_{\text{filt}}(T^-, T^+)$. The proportions $N_{\text{type}}(T)/N_{\text{filt}}(T^-, 2)$, which is the main quantity of interest, can be obtained at any $T \in [-6, 2]$ if $T^- = -16$, and at just one value of $T = 2$ if $T^- = -8$. Therefore, to make a meaningful comparison of CT classifications for different values of T^-, we focus on the classifications at $T = 2$, that is two seconds after the time at which a CT is detected. In other words, we examine $N_{\text{type}}(2)/N_{\text{filt}}(T^-, 2)$ for $T^- = -16, -14, -12, -10$, and -8, where type = BCT, BNCT or NCT is given by the SVM classifier.

Key findings from the main experiment

The key findings from our main experiment are presented in the two rows in Fig. 9, bearing in mind that the model-trained SVM uses the eight features obtained from the portion of the voltage recording that starts $|T^-|$ seconds prior to and always ends $T^+ = 2$ seconds after the detection of each filtered CT. The results presented in Fig. 9 are based on the classifications obtained at time $T = T^+ = 2$.

The top row of Fig. 9 shows how the proportions of all filtered CTs that are classified as (red) NCTs, (blue) BCTs, and (green) BNCTs vary with T^-. We show this for rat S in panel (a), rat T in panel (b) and rat K in panel (c). The key finding is that NCTs are the dominant seizure generation mechanism in all three rats for all values of T^-. More specifically,

- NCTs account for approximately $41 - 51\%$ of the filtered CTs,
- BCTs account for approximately $22 - 34\%$ of the filtered CTs, and
- BNCTs account for approximately $22 - 27\%$ of the filtered CTs.

In the bottom row of Fig. 9, we show how the proportion of filtered CTs to the total number of detected CTs increases with T^- for a fixed $T^+ = 2$. For the minimum value of T^-, which is -8 and is a technical restriction of our classifier, approximately 50% of the detected CTs are retained by filtering (i.e., are classifiable) for rat S in panel (d), compared to approximately 40% for rat T in panel (e), and approximately 35% for rat K in panel (f) partially due to many artefacts. Finally, a numerical summary of the key findings is provided in Table 4.

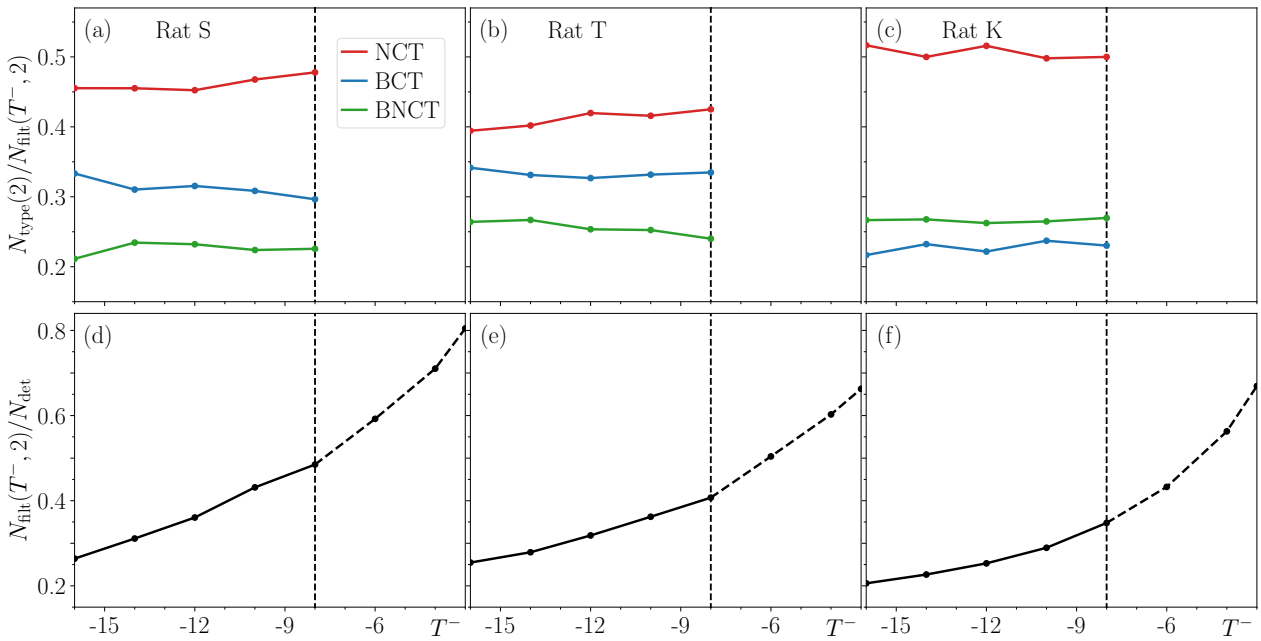

Figure 9. (Top row) The proportion of filtered CTs from the voltage recordings in GAERS that are classified by the model-trained SVM, i.e. SVM type-3 with $t_w = 1$ and $t_m = 8$, at time $T = T^+ = 2$ as (red) NCT, (blue) BCT and (green) BNCT vs. T^- for (a) rat S, (b) rat T and (c) rat K. Note that NCT emerges as the dominant CT-type responsible for the onset of seizures in all three rats. (Bottom row) The proportion of filtered CTs to the total number of CTs detected in the voltage recordings vs. T^- for (d) rat S, (e) rat T and (f) rat K. Other parameters are specified in Table 3, where α is chosen according to the method outlined in S5 in [30].

The bottom row in Fig. 9 shows that many detected CTs were excluded from classification due to filtering, i.e., they did not meet the technical requirements of our classifier, namely conditions C1-C5 with $T^- \leq -8$. This raises the question of how the exclusion of these detected CTs impacted the key findings. To address this question, we make three observations. First, if T^- were increased to -3, over 80% of the detected CTs for rat S would meet the filtering criteria C1-C5. Second, the top row in Fig. 9 shows that the proportion of classified CTs of a given type in all three rats remains largely unchanged as T^- is varied from -16 to -8. This suggests that the trend continues for larger values of T^-, but this cannot be verified due to the technical requirements of our classifier. Third, McCafferty *et al.*[22] found that approximately 50% of seizures in GAERS can be interrupted by applying a suitably chosen auditory stimulus. This suggests that, like in our dominant case of noise-induced CT, and the least dominant case of bifurcation/noise-induced CT, epileptic brain frequently operates in a multistable regime whereby disturbances can force transitions between coexisting stable NS and S states. For these reasons, we conjecture that the key findings persist for most of the excluded detected CTs.

Discussion

This paper presents a general framework for classifying seizure generation mechanisms as different types of critical transition (CT) between the non-seizure state (NS state) and the seizure state (S state) in the epileptic brain. The framework consists of four steps. The first step is to construct a canonical mathematical model which displays CTs that (i) closely resemble voltage recordings of real seizures and (ii) can be bifurcation-induced (BCTs), noise-induced (NCTs), or a combination of both (BNCTs). The second step is to construct a detection algorithm that detects CTs between the NS and S states in both the model's output and voltage recordings of real seizures. The third step is to construct a machine learning CT-type classifier. The classifier is trained using selected time series properties of the model's output where the CT-types are known. In the fourth step, we use the model-trained classifier on voltage recordings, where CT-types are unknown, in order to classify the onset of real seizures as BCTs, NCTs or BNCTs.

We applied the above framework to voltage recordings taken from inside the brain of *Genetic Absence Epilepsy Rats from Strasbourg* (GAERS) - a well-established model of absence epilepsy - and, for the first time, demonstrated that consensus can

	Rat S	Rat T	Rat K
No. of CTs detected by expert	621	1593	1136
No. of CTs detected by the algorithm (N_{det})	466	1115	874
No. of filtered CTs ($N_{\text{filt}}(-8,2)$)	226	454	304
Percentage of filtered CTs classified as type = BCTs, BNCTs, or NCTs ($(N_{\text{type}}(2)/N_{\text{filt}}(-8,2))\times 100$)	29.6, 22.6, **47.8**	33.5, 24.0, **42.5**	23.0, 27.0, **50.0**

Table 4. Our key findings in numbers. In the bottom row, the dominant CT-type for each rat is highlighted by the bold font.

be reached on which CT-types are predominantly responsible for seizure onset. We found NCTs to be the dominant seizure generation mechanism - they accounted for almost 50% of the seizures we classified, with BCTs and BNCTs accounting for the remainder. Our findings show that (i) seizures in the brain appear to be generated by different CT-types[40–45], and (ii) multistability and noise play a significant role in the epileptic brain[16]. They also challenge the view that seizures are predominantly bifurcation-induced[15].

An important practical implication of our findings is that different CT-types may reflect distinct neural processes involved in generating seizures. From this perspective, a seizure that is classified as an NCT is likely to start locally with respect to a recording electrode. Furthermore, there are constitutively pro-ictal neurons in the epileptic brain that can initiate a seizure based solely on noisy input in the form of exogenous or endogenous disturbances to these neurons. The hyper-excitable neurons found in layer V of the perioral somatosensory cortex[46], from which GAERS seizures appear to originate[47], may be one such example. On the other hand, a seizure that is classified as a BCT is likely to emerge gradually, or propagate to the recording electrode having started elsewhere. In this case, the slow drift in brain state is considered as the ictogenic culprit. The observed changes in brain state tens of seconds prior to seizure onset in rats and humans[48,49] may reflect such a gradual drift towards seizure onset. Therefore, identifying and arresting this slow drift may prevent seizure onsets. BNCTs may represent combinations of the neural processes discussed above.

Classifying seizure generation mechanisms in a given brain is not only an important fundamental question. It could transform how we model, detect, prevent, and treat seizures, with the ultimate aim of informing new clinical decision support systems and tailored seizure prevention[45] and control[50] strategies. While studies in GAERS have identified specific neurons capable of seizure initiation, human epilepsy may have multiple potential cortical sites of origin, likely related to cortical-thalamic interactions[19]. Furthermore, these sites may differ between babies and adults. It remains to be seen whether different CT-types correspond to how people with epilepsy experience their seizures, or whether NCTs play a dominant role. Our framework provides proof of concept and represents an important first step towards analysing those more complex seizures. The next step is to apply it to the less well-understood forms of human epilepsy. This may involve updating our mathematical model of seizure activity with a more realistic one, augmenting our CT detection algorithm with additional time series properties such as the dynamical eigenvalue early warning signal proposed by Grziwotz *et al.*[51], incorporating different experimental data such as heart rate variability[52], and developing more advanced filtering techniques to remove artefacts[53]. Furthermore, our framework can be extended to include other types of CT, such as rate-induced CTs[54], identify new forms of early warning signals for different CT-types, and classify seizure termination mechanisms[55,56] to gain new insight into how the brain changes in response to a seizure and how an ongoing seizure can be terminated. Finally, it can also be applied to complex systems beyond the brain.

Methods

M1: Obtaining and annotating seizure data from GAERS

The data discussed here was obtained by Cian McCafferty, François David, and Vincenzo Crunelli and was presented in 18.

Obtaining the data: The electrophysiological data was acquired from male *Genetic Absence Epilepsy Rats from Strasbourg* (GAERS) aged between 4-7 months when in a state of relaxed wakefulness where seizures were experienced more often. Silicon-site electrodes were used to sample voltages (20000/second) from the *ventrobasal thalamus* while the rats were able to move freely and alternate between waking, sleeping, and seizing states. This data was processed with a Plexon HST/32V-G20 VLSI-based preamplifier and associated digitization system and subsequently down-sampled to 1000 samples/second.
Labelling the recording session: Each recording session was labelling using the following convention described through example, 'S1K': voltage recordings from rat S on day 1 of recording during the Kth recording session on that day.
Annotating the data: Spike-wave discharges (SWDs) were identified using Cambridge Electronic Design's software, 'Spike2'. In all cases, EEG at 1000Hz was used for this step. EEG was acquired at 1000Hz for fMRI (see McCafferty *et al.*[48]) and behaviour was down-sampled by averaging for neuronal activity. Briefly, smoothening (voltage at time t is set to mean of voltages from time $t-10$ms to time $t+10$ms) and DC removal (voltage at time t is set to original voltage minus mean of voltages from time $t-0.1$s to time $t+0.11$s) functions were used to reversibly visually clean the frontoparietal differential EEG. Then, a negative amplitude threshold (mean voltage minus 5-7 standard deviations of baseline non-SWD EEG) was used to detect putative spike-wave crossing points, defined as whenever the signal crossed this amplitude threshold. The crossing points were then grouped into events based on the intervals between them (maximum time between initial two crossings 0.2s, maximum time between any two crossings within an event 0.35s, minimum of 5 crossings per event) and the defined properties of SWDs (minimum duration 0.5s, minimum inter-SWD interval 0.5s merging any SWDs with shorter intervals), and subsequently, using a frequency threshold, these events were classified as SWDs (if > 75% of intercrossing intervals were within a 5–12Hz range) or other (e.g., noise, sleep). Labelled SWDs were then visually inspected for accuracy of SWD detection, as well as onset and offset times. Periods of sleep were identified based on sharp increases in the 1–4Hz frequency band and were excluded from analysis. Periods of non-REM sleep were rare in the recordings due to their relatively short duration and therefore not informative for analysis.
Artefacts and how they are accounted for: Artefacts in the voltage recordings can appear for different reasons, the most common being (i) movement of electrodes/wiring and (ii) movement of muscles close to the electrodes (generally for chewing). Some artefacts are generated by signal overload and noise from electrical devices around the recording setup, however, these are less common as the recording process was generally appropriately amplified and shielded. Artefacts are accounted for through the method described above since high-amplitude events that do not align with the frequency profile of SWDs were not annotated. Additionally, the visual inspection described above also includes the manual rejection of artefacts which could be excised from recordings. Entire recording session were excluded if more than 5% of the session consisted of artefact activity.

Remark (comparison with human data): GAERS can express multiple absence seizures per minute[57], far exceeding the frequency of any absence seizure syndrome in humans. For instance, Gregorčič *et al.*[58] found that, for children with treatment-resistant childhood absence epilepsy, the median number of seizures per day was three. For a review of the GAERS model, with attention to its similarities and differences to human absence seizures, see Depaulis *et al.*[59].

M2: Computing the time series properties and their slopes

Summary of procedures used to compute different properties of time series of both the model's output and the voltage recordings.
Detrending the time series: We first detrend a given time series using the Gaussian filtering approach mentioned in Dakos et al.[60]. This involves fitting a Gaussian kernel smoothing function to the time series and then subtracting the fit from the time series to obtain the detrended time series. As further detailed in Lenton *et al.*[61], this process involves choosing a suitable bandwidth (a fitting parameter) for the kernel which determines the degree of smoothing. The bandwidth is typically chosen to neither over-fit the data nor filter out low frequencies from the time series, in our case we choose a bandwidth of 30. In our experiments we use the 'gaussian_filter1d' function from the 'scipy.ndimage' Python library to detrend the time series. See Fig. 10 for an example of a detrended time series of voltage recordings. In our experiments we find that, whether the data is detrended or not, there are minor changes in the main results presented in Figs. 9 (a)-(c), these changes are on the order of 2-3% in terms of the percentage of CTs classified as BCTs, BNCTs, and NCTs specified in Table 4. Note, we only detrend voltage recordings as no trends are present in our model's output.
Variance: We compute the Gaussian rolling variance at time $t \geq t_w$ using data points of the time series within the interval, $[t-t_w, t]$, which we call 'the window' and $t_w > 0$ is the 'window length'. The variance is calculated as the weighted variance of a Gaussian distribution of the data points centered at $t-t_w/2$ with mean equal to zero. This involves assigning weights to data points within the window and these weights are determined from the probability density function of a Gaussian distribution

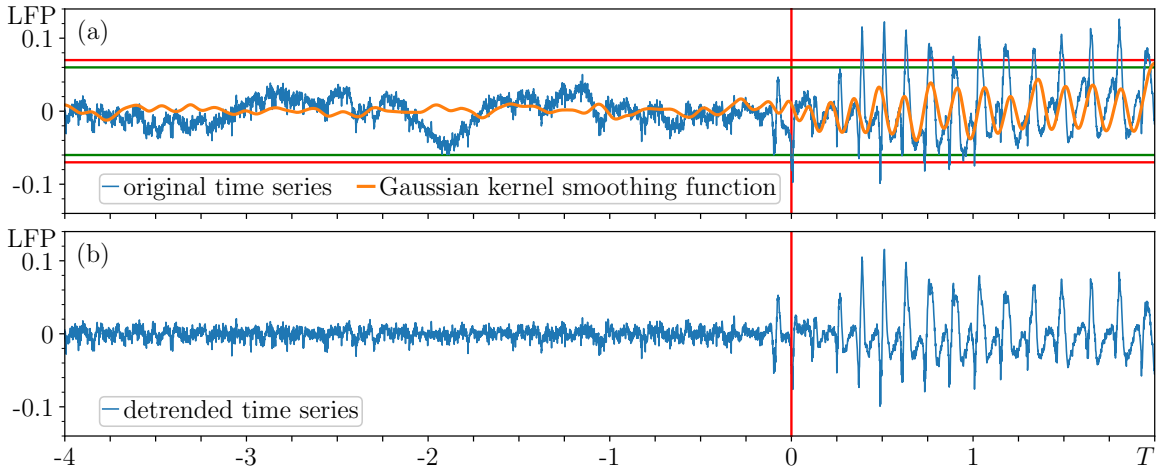

Figure 10. Detrending voltage recordings using a Gaussian kernel smoothing function. (a) shows (in blue) an example of voltage recordings of real seizure activity before detrending and (in orange) the corresponding function fitted to these voltage recordings, and (b) shows the detrended voltage recordings. The vertical red solid line indicates when $T = 0$. Horizontal lines indicate the thresholds of (in red) α and (in green) β used by the algorithm. The other algorithm parameters are specified in Table 3.

centered at the middle of the window with standard deviation taken as $t_w/6$. Based on this procedure we define $\mathrm{GV}(t) : \mathbb{R} \to \mathbb{R}$ as the Gaussian variance of the time series at $t \geq t_w$.

Logarithm of the variance: Considering the results presented by Milanowski and Suffczynski[12], we also find it useful to record the base-10 logarithm of the GV in each of the moving windows described above. This allows us to more closely examine the smaller changes in the variance. We define $\log_{10}\mathrm{GV}(t) : \mathbb{R} \to \mathbb{R}$ as the base-10 logarithm of the Gaussian variance of the time series at $t \geq t_w$.

Autocorrelation: We compute the autocorrelation at $t \geq t_w$ using a rolling window technique and data points of the time series within the same intervals as above, $[t - t_w, t]$. More specifically, we compute the lag-1 autocorrelation at a given time t by computing the Pearson correlation between data points in the following two intervals, $[t - t_w + t_a, t]$ and $[t - t_w, t - t_a]$ where $t_a > 0$ is the time between two successive points in the time series. Based on this procedure we define $\mathrm{AC}(t) : \mathbb{R} \to \mathbb{R}$ as the lag-1 autocorrelation of the time series at $t \geq t_w$.

Variance of the autocorrelation: Prior to a BCT or BNCT from the NS to S state we observe a decrease in the fluctuations of $\mathrm{AC}(t)$ and, after the CT, $\mathrm{AC}(t)$ remains relatively constant and close to 1. To more closely examine this behaviour, we compute the base-10 logarithm of the Gaussian variance of the time series obtained for $\mathrm{AC}(t)$ using the same procedure outlined above. From this we define $\log_{10}\mathrm{GV}(\mathrm{AC})(t) : \mathbb{R} \to \mathbb{R}$ as the base-10 logarithm of the Gaussian variance of the autocorrelation of a given time series at $t \geq 2\,t_w$ (since $\mathrm{AC}(t)$ is defined for $t \geq t_w$, the resulting $\log_{10}\mathrm{GV}(\mathrm{AC})(t)$ is defined for $t \geq 2\,t_w$).

For convenience we refer to $\mathrm{GV}(t)$ as GV, and similarly for $\log_{10}$GV, AC, and $\log_{10}$GV(AC). We collectively refer to these four quantities as the time series properties (TSPs). From several trial runs we find that setting $t_w = 1$ provides relatively robust estimates of the variance and autocorrelation of the time series. Each window of length t_w is shifted forward in time by a factor $\Delta_w > 0$ which we choose as 0.001, the time between two successive points in the original time series.

Slopes: We also analyse how the slopes of the TSPs change over time by fitting the function, $f(x) = m\,x + c$, to the values of, for example, GV within the interval $[t - t_m, t]$, where t_m is the 'slope length'. We use the 'curve_fit()' function in the SciPy Python library to perform a least squares fit of this function to the data via the Levenberg–Marquardt algorithm. Based on this procedure we define $m(\mathrm{GV}, t_m) : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ as the slope of the Gaussian variance of time series at $t \geq t_w + t_m$. The same procedure is applied to all TSPs. Note, $m(\log_{10}\mathrm{GV(AC)},\ t_m)$ is defined for $t \geq 2\,t_w + t_m$. We collectively refer to these four quantities as 'the slopes of the TSPs for a given t_m', denoted by $m(\mathrm{TSPs},\ t_m)$. Each window of length t_m is shifted forward in time by a factor $\Delta_m > 0$ which we choose as $\Delta_m = 100\Delta_w$ to reduce computational time.

In Figs. 11 and 12 we illustrate the different time intervals over which TSPs and $m(\mathrm{TSPs},\ t_m)$ are computed using an example of the model's output nearby a BCT from the NS to S state. This helps clarify (i) how values of TSPs and m(TSPs,

	Gaussian variance	base-10 logarithm of the Gaussian variance	lag-1 autocorrelation	base-10 logarithm of the Gaussian variance of the lag-1 autocorrelation
Time series properties (TSPs)	GV	$\log_{10}$GV	AC	$\log_{10}$GV(AC)
Slopes of time series properties (m(TSPs, t_m))	m(GV, t_m)	m($\log_{10}$GV, t_m)	m(AC, t_m)	m($\log_{10}$GV(AC), t_m)

Table 5. Quantities used as features to classify seizure generation mechanisms in a noisy time series. These quantities are dependent on time, i.e., GV = GV(t), and are computed using data points of a given time series within the interval $[t-t_w, t]$ where $t_w > 0$ is the 'window length'. t_m denotes the slope length, the duration of time over which the slope is calculated.

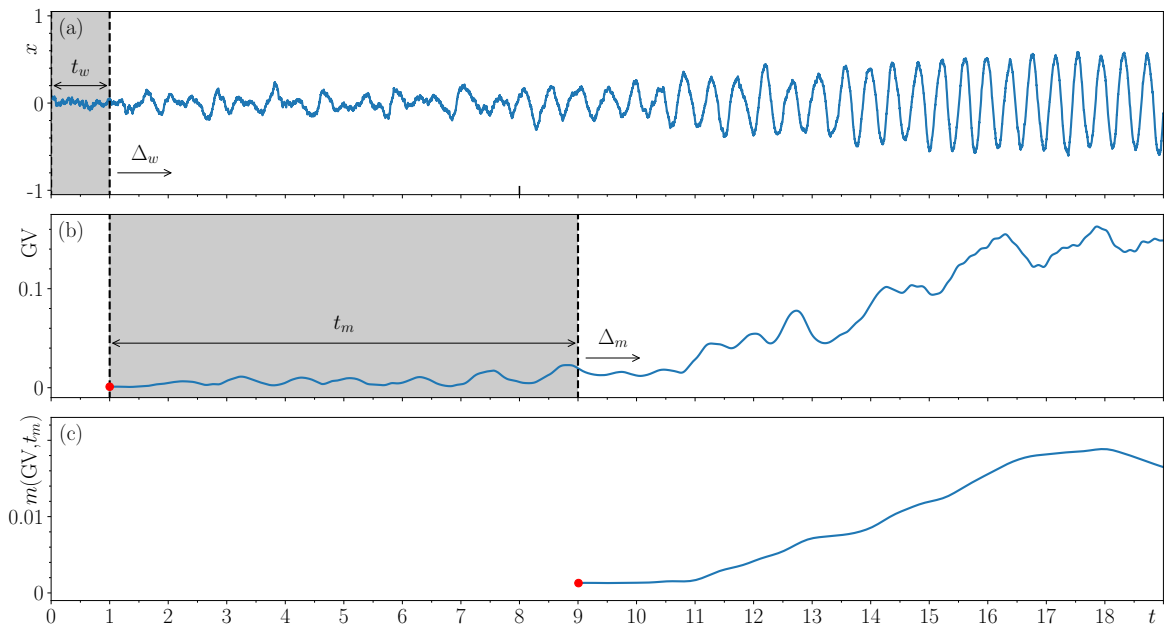

Figure 11. Illustrating process for computing values of GV and m(GV, t_m) using moving windows of lengths t_w and t_m, which are advanced along the respective time series in steps of Δ_w and Δ_m. The shaded region shows values of: (a) the model's output used to compute (red point) the first GV value in (b) and (b) GV used to compute (red point) the first m(GV, t_m) value in (c). The time series in (a) is an example of the model's output around a BCT from the NS to S state at $t = 8$ (inward black tick), generated using parameters specified in Table 2 and plotted here for a new time $t = t' + 38$, where t' is the original time, to better illustrate the above process. $t_w = 1$, $\Delta_w = 0.001$, $t_m = 8$, and $\Delta_m = 0.1$.

t_m) are obtained from moving windows and (ii) the correspondence in time between TSPs, m(TSPs, t_m), and the original time series from which they are obtained. More specifically, Fig. 11 illustrates the above for GV and m(GV, t_m), the same procedure is used to compute AC, $\log_{10}$GV, and their corresponding slopes. Fig. 12 illustrates the above for $\log_{10}$GV(AC) and m($\log_{10}$GV(AC), t_m). In both Figs. 11 and 12, $t_w = 1$ and $t_m = 8$.

M3: SVM training and testing

We use the 'sklearn.svm' Python library to construct a linear support vector machine (SVM) that classifies CTs from the NS to S state according to the CT-type where type = BCT, BNCT, or NCT. The SVM classifier is obtained by finding an optimal line/hyperplane that maximises the distance between each class, i.e., CT-type, in an F-dimensional space where $F \in \mathbb{N}^+$ is the number of 'features'.

Preparing the data: We use values of the TSPs and m(TSPs, t_m) at a time t as features. $S \in \mathbb{N}$ samples of each feature are used to train the SVM. Since (i) some features are orders of magnitude larger than others and (ii) we would like to construct a classifier which treats features with equal importance, each of the features are scaled to be of zero mean and unit variance using the 'preprocessing.scale' function from the 'sklearn' Python library. The data used to construct the SVM classifier is structured as follows, the 'feature matrix' is defined as $X \in \mathbb{R}^{S\times F}$, and the 'class vector' is defined as $y \in \mathbb{R}^S$. The information contained

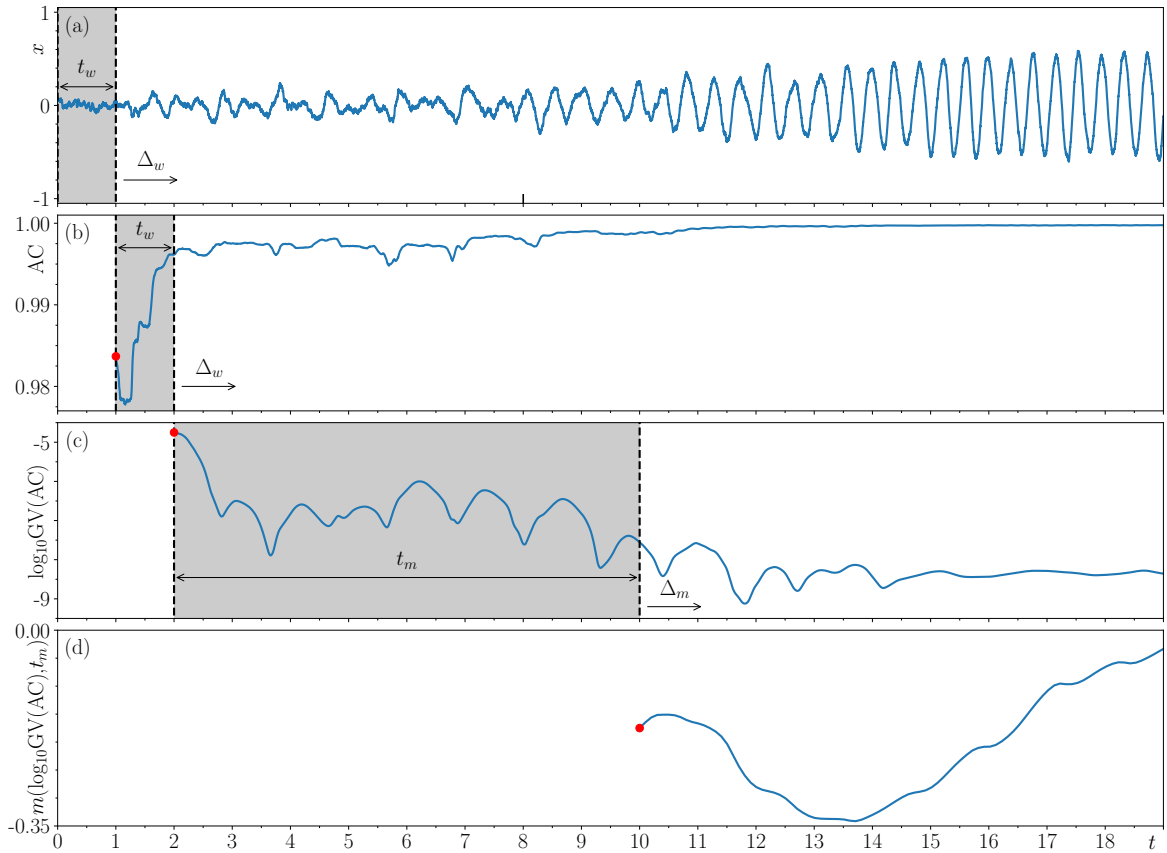

Figure 12. Illustrating process for computing values of $\log_{10}$GV(AC) and $m(\log_{10}\text{GV(AC)}, t_m)$ using moving windows of lengths t_w and t_m, which are advanced along the respective time series in steps of Δ_w and Δ_m. The shaded regions indicate values of: (a) the model's output used to compute (red point) the first AV value in (b), (b) AC used to compute (red point) the first $\log_{10}$GV(AC) value in (c), and (c) $\log_{10}$GV(AC) used to compute (red point) the first $m(\log_{10}\text{GV(AC)}, t_m)$ value in (d). The time series in (a) is the same as Fig. 11 (a) (see caption for details), $t_w = 1$, $\Delta_w = 0.001$, $t_m = 8$, and $\Delta_m = 0.1$.

in X and y define our data set.

Training and testing: We split the above data set into a training and testing data sets whereby 70% of the data is used for training and 30% for testing (a convention commonly followed when training and testing SVMs.) The SVM classifier is constructed based on the training data set and the accuracy of the SVM classifier is determined based on the fraction of correctly classified samples in the testing data set.

Feature importance: We use the 'permutation_importance' function from the 'sklearn.inspection' Python library to determine what features have the greatest impact on the accuracy. This function constructs different random permutations of a given feature (i.e., reorders elements in a column of X reserved for testing while the corresponding y remains unchanged) and returns the resulting decrease in accuracy; the greater the decrease the more dependent the SVM classifier's accuracy is on the feature. This process is repeated 100 times for each feature to obtain an estimate of the 'mean permutation importance' (MPI) where $-1 \leq \text{MPI} \leq 1$. We also obtain the corresponding standard deviations from this process. Therefore, if, for a given feature,

- MPI > 0, the SVM classifier's accuracy depends on this feature and the greater the MPI the more important the feature.
- MPI ≈ 0, the SVM classifier's accuracy does not depend on this feature.
- MPI < 0, the SVM classifier's accuracy increases when using the permuted data set as opposed to the original data set.

	Features used
SVM type-1	TSPs: GV, $\log_{10}$GV, AC, and $\log_{10}$GV(AC).
SVM type-2	m(TSPs, t_m) for a given t_m: m(GV, t_m), m($\log_{10}$GV, t_m), m(AC, t_m), and m($\log_{10}$GV(AC), t_m).
SVM type-3	All TSPs and m(TSPs, t_m) for a given t_m.

Table 6. Features used to train and test different types of Support Vector Machines (SVMs). The time series properties (TSPs) and their slopes, m(TSPs, t_m), are specified in Table 5. t_m denotes the duration of time over which the slope is calculated.

References

1. Beniczky, S. *et al.* Updated classification of epileptic seizures: Position paper of the International League Against Epilepsy. *Epilepsia* **66**, 1804–1823 (2025).
2. Kaye, D. *et al.* Impact of Prolonged Seizures on Patients' and Caregivers' Quality of Life (P1-9.014). In *Neurology*, vol. 104, 2708 (Lippincott Williams & Wilkins Hagerstown, MD, 2025).
3. González, O. C., Krishnan, G. P., Timofeev, I. & Bazhenov, M. Ionic and synaptic mechanisms of seizure generation and epileptogenesis. *Neurobiol. Dis* **130**, 104485 (2019).
4. Ashwin, P., Perryman, C. & Wieczorek, S. Parameter shifts for nonautonomous systems in low dimension: bifurcation-and rate-induced tipping. *Nonlinearity* **30**, 2185 (2017).
5. Lenton, T. M. *et al.* Tipping elements in the Earth's climate system. *Proc. Natl. Acad. Sci.* **105**, 1786–1793 (2008).
6. Lenton, T. M. Tipping positive change. *Philos. Trans. R. Soc. B* **375**, 20190123 (2020).
7. Armstrong McKay, D. I. *et al.* Exceeding 1.5 C global warming could trigger multiple climate tipping points. *Science* **377**, eabn7950 (2022).
8. Lenton, T. M. *et al.* Global tipping points report 2025 (2025). University of Exeter.
9. Scheffer, M. *et al.* Early-warning signals for critical transitions. *Nature* **461**, 53 (2009).
10. McSharry, P. E., Smith, L. A. & Tarassenko, L. Prediction of epileptic seizures: are nonlinear methods relevant? *Nat. Med.* **9**, 241–242 (2003).
11. Meisel, C. & Kuehn, C. Scaling effects and spatio-temporal multilevel dynamics in epileptic seizures. *PLoS One* **7**, e30371 (2012).
12. Milanowski, P. & Suffczynski, P. Seizures start without common signatures of critical transition. *Int. J. Neural Syst.* **26**, 1650053 (2016).
13. Jirsa, V. K., Stacey, W. C., Quilichini, P. P., Ivanov, A. I. & Bernard, C. On the nature of seizure dynamics. *Brain* **137**, 2210–2230 (2014).
14. Wilkat, T., Rings, T. & Lehnertz, K. No evidence for critical slowing down prior to human epileptic seizures. *Chaos* **29**, 091104 (2019).
15. Maturana, M. I. *et al.* Critical slowing down as a biomarker for seizure susceptibility. *Nat. Commun.* **11**, 2172 (2020).
16. da Silva, F. L. *et al.* Epilepsies as Dynamical Diseases of Brain Systems: Basic Models of the Transition Between Normal and Epileptic Activity. *Epilepsia* **44**, 72–83 (2003).
17. Dakos, V. *et al.* Tipping point detection and early warnings in climate, ecological, and human systems. *Earth Syst. Dyn.* **15**, 1117–1135 (2024).
18. McCafferty, C. *et al.* Cortical drive and thalamic feed-forward inhibition control thalamic output synchrony during absence seizures. *Nat. neuroscience* **21**, 744–756 (2018).
19. Crunelli, V. *et al.* Clinical and experimental insight into pathophysiology, comorbidity and therapy of absence seizures. *Brain* **143**, 2341–2368 (2020).
20. Zhou, Z. *et al.* A generalized seizure type: Myoclonic-to-tonic seizure. *Clin. Neurophysiol.* **164**, 24–29 (2024).
21. Depaulis, A., van Luijtelaar, G., Pitkanen, A., Schwartzkroin, P. & Moshe, S. Models of seizures and epilepsy (2006).
22. McCafferty, C. P., Zheng, X., Tung, R., Gruenbaum, B. F. & Blumenfeld, H. Interruption of rat absence seizures by auditory stimulation. *bioRxiv* 2025–10 (2025).

23. da Silva, F. H. L. *et al.* Dynamical diseases of brain systems: different routes to epileptic seizures. *IEEE Trans. Biomed. Eng.* **50**, 540–548 (2003).

24. Suffczynski, P., Kalitzin, S. & da Silva, F. L. Dynamics of non-convulsive epileptic phenomena modeled by a bistable neuronal network. *Neuroscience* **126**, 467–484 (2004).

25. Suffczynski, P., da Silva, F. L., Parra, J., Velis, D. & Kalitzin, S. Epileptic transitions: model predictions and experimental validation. *J. Clin. Neurophysiol.* **22**, 288–299 (2005).

26. Junges, L., Woldman, W., Benjamin, O. J. & Terry, J. R. Epilepsy surgery: Evaluating robustness using dynamic network models. *Chaos* **30**, 113106 (2020).

27. Harrington, E. G., Kissack, P., Terry, J. R., Woldman, W. & Junges, L. Treatment effects in epilepsy: a mathematical framework for understanding response over time. *Front. Netw. Physiol.* **4**, 1308501 (2024).

28. Qin, Y., El-Gazzar, A., Bassett, D. S., Pasqualetti, F. & Van Gerven, M. Analytical characterization of epileptic dynamics in a bistable system. In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 583–588 (IEEE, 2024).

29. Byrne, Á., Ross, J., Nicks, R. & Coombes, S. Mean-field models for EEG/MEG: from oscillations to waves. *Brain Topogr.* **35**, 36–53 (2022).

30. Flynn, A. *et al.* Supplementary material from: "Classifying seizure generation mechanisms: A critical transitions framework". (Link).

31. Kuznetsov, Y. A. *Elements of Applied Bifurcation Theory* (Springer Science & Business Media, 2013).

32. Lin, K. K. & Young, L.-S. Shear-induced chaos. *Nonlinearity* **21**, 899 (2008).

33. Wieczorek, S. Stochastic bifurcation in noise-driven lasers and Hopf oscillators. *Phys. Rev. E* **79**, 036209 (2009).

34. Blackbeard, N., Erzgräber, H. & Wieczorek, S. Shear-induced bifurcations and chaos in models of three coupled lasers. *SIAM J. on Appl. Dyn. Syst.* **10**, 469–509 (2011).

35. Krasnosel'skii, M. A. & Pokrovskii, A. V. *Systems with Hysteresis* (Springer Science & Business Media, 2012).

36. Ditlevsen, P. & Ditlevsen, S. Warning of a forthcoming collapse of the Atlantic meridional overturning circulation. *Nat. Commun.* **14**, 4254 (2023).

37. Van Westen, R. M., Kliphuis, M. & Dijkstra, H. A. Physics-based early warning signal shows that AMOC is on tipping course. *Sci. Adv.* **10**, 1189 (2024).

38. Ashwin, P., Bastiaansen, R., von der Heydt, A. S. & Ritchie, P. D. Early warning skill, extrapolation and tipping for accelerating cascades. *Proc. Roy. Soc. A* **481**, 20250405 (2025).

39. Bury, T. M. *et al.* Deep learning for early warning signals of tipping points. *Proc. Natl. Acad. Sci.* **118**, e2106140118 (2021).

40. McCormick, D. A. & Contreras, D. On the cellular and network bases of epileptic seizures. *Annu. Rev. Physiol.* **63**, 815–846 (2001).

41. Vezzani, A., French, J., Bartfai, T. & Baram, T. Z. The role of inflammation in epilepsy. *Nat. Rev. Neurol.* **7**, 31–40 (2011).

42. Kuhlmann, L., Lehnertz, K., Richardson, M. P., Schelter, B. & Zaveri, H. P. Seizure prediction — ready for a new era. *Nat. Rev. Neurol.* **14**, 618–630 (2018).

43. González, O. C., Krishnan, G. P., Timofeev, I. & Bazhenov, M. Ionic and synaptic mechanisms of seizure generation and epileptogenesis. *Neurobiol. Dis.* **130**, 104485 (2019).

44. Jiruska, P., Freestone, D., Gnatkovsky, V. & Wang, Y. An update on the seizures beget seizures theory. *Epilepsia* **64**, S13–S24 (2023).

45. Lehnertz, K., Broehl, T. & von Wrede, R. Epileptic-network-based prediction and control of seizures in humans. *Neurobiol. Dis.* **181**, 106098 (2023).

46. Polack, P.-O. *et al.* Deep layer somatosensory cortical neurons initiate spike-and-wave discharges in a genetic model of absence seizures. *J. Neurosci.* **27**, 6590–6599 (2007).

47. Meeren, H. K., Pijn, J. P. M., Van Luijtelaar, E. L., Coenen, A. M. & da Silva, F. H. L. Cortical focus drives widespread corticothalamic networks during spontaneous absence seizures in rats. *J. Neurosci.* **22**, 1480–1495 (2002).

48. McCafferty, C. *et al.* Decreased but diverse activity of cortical and thalamic neurons in consciousness-impairing rodent absence seizures. *Nat. Commun.* **14**, 117 (2023).

49. Bai, X. *et al.* Dynamic time course of typical childhood absence seizures: EEG, behavior, and functional magnetic resonance imaging. *J. Neurosci.* **30**, 5884–5893 (2010).

50. Gluckman, B. J., Nguyen, H., Weinstein, S. L. & Schiff, S. J. Adaptive electric field control of epileptic seizures. *J. Neurosci.* **21**, 590–600 (2001).

51. Grziwotz, F. *et al.* Anticipating the occurrence and type of critical transitions. *Sci. Adv.* **9**, eabq4558 (2023).

52. Rezaei, K. *et al.* Assessing the effectiveness of heart rate variability as a diagnostic tool for brain injuries in infants. In *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 1–4 (2024).

53. Zhang, X., Zhang, X., Huang, Q. & Chen, F. A review of epilepsy detection and prediction methods based on EEG signal processing and deep learning. *Front. Neurosci.* **18**, 1468967 (2024).

54. Ritchie, P. D., Alkhayuon, H., Cox, P. M. & Wieczorek, S. Rate-induced tipping in natural and human systems. *Earth Syst. Dyn.* **14**, 669–683 (2023).

55. Kramer, M. A. *et al.* Human seizures self-terminate across spatial scales via a critical transition. *Proc. Natl. Acad. Sci.* **109**, 21116–21121 (2012).

56. Schindler, K., Leung, H., Elger, C. E. & Lehnertz, K. Assessing seizure dynamics by analysing the correlation structure of multichannel intracranial EEG. *Brain* **130**, 65–77 (2007).

57. Powell, K. L. *et al.* Seizure expression, behavior, and brain morphology differences in colonies of Genetic Absence Epilepsy Rats from Strasbourg. *Epilepsia* **55**, 1959–1968 (2014).

58. Gregorčič, S. *et al.* Difficult to treat absence seizures in children: A single-center retrospective study. *Front. Neurol.* **13**, 958369 (2022).

59. Depaulis, A., David, O. & Charpier, S. The genetic absence epilepsy rat from strasbourg as a model to decipher the neuronal and network mechanisms of generalized idiopathic epilepsies. *J. Neurosci. Methods* **260**, 159–174 (2016).

60. Dakos, V. *et al.* Slowing down as an early warning signal for abrupt climate change. *Proc. Natl. Acad. Sci.* **105**, 14308–14312 (2008).

61. Lenton, T. M., Livina, V., Dakos, V., van Nes, E. H. & Scheffer, M. Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness. *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.* **370**, 1185–1204 (2012).

Acknowledgements

This publication has emanated from research conducted with the financial support of Taighde Éireann – Research Ireland under grant number [19/FFP/6782].

Author contributions statement

A.F. and S.W. conceived the experiments, A.F. conducted the experiments, A.F. and S.W. analysed the results, C.Mc C., F.D., and V.C. provided the voltage recordings, C.Mc C., K.L., and W.P.M. provided valuable discussions. A.F., C.Mc C., K.L., W.P.M., and S.W. reviewed the manuscript.

Additional information

The authors have no competing interests to disclose.

Supplementary information

S1: Applying the algorithm to detect seizure-like activity in the model's output

In this subsection we illustrate some common outcomes that arise when using our algorithm to detect CTs in noisy time series of the model's output for each CT-type it can produce, namely BCTs, BNCTs, and NCTs.

NCTs: We first integrate the model from $t = 0$ to $t = 200000$ with model parameters specified in Table 2. We then apply our algorithm to (x) the model's output using the parameter values specified in Table 3. The algorithm detected (i) 193 CTs between the NS and S states (i.e. 193 pairs of t_1 and t_2 values) and (ii) 86 almost-occurring CTs (i.e., a set of 86 values consisting of different $\tilde{t}_1$ and $\tilde{t}_2$ values), examples of these are shown in Fig. S-1. For our convenience, in Figs. S-1 (a)-(d) we plot x versus $T = t - t_1$ for the chosen values of t_1, and in Fig. S-1 (e) we plot x versus $\tilde{T} = t - \tilde{t}_1$ for the chosen value of $\tilde{t}_1$. Solid vertical lines indicate (in red) when $T = 0$ and (in green) when $T = t_2 - t_1$. Similarly, dashed vertical lines indicate (in red) $\tilde{T} = 0$ and (in green) $\tilde{T} = \tilde{t}_2 - \tilde{t}_1$. The horizontal lines indicate the values of (in red) α and (in green) β used by the algorithm.

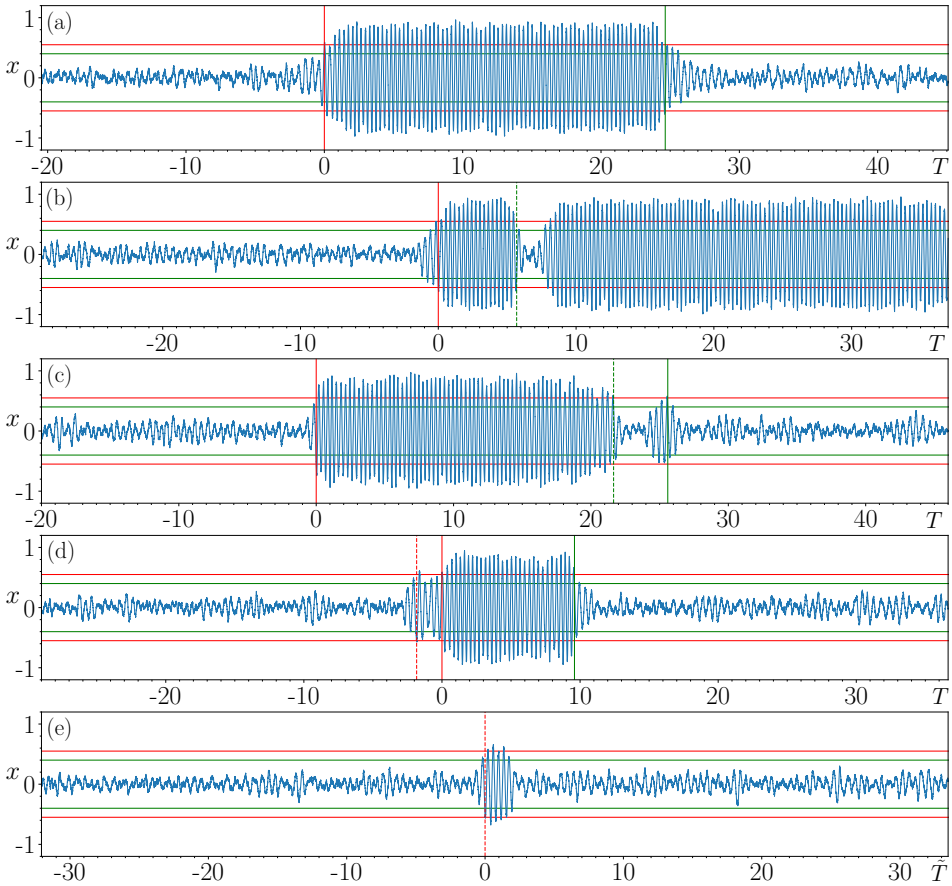

Figure S-1. Examples of NCTs in the model's output (the variable x from System (5)) that were detected by the algorithm. The model's output is plotted versus $T = t - t_1$ for chosen values of t_1 in (a)-(d), and versus $\tilde{T} = t - \tilde{t}_1$ for a chosen $\tilde{t}_1$ in (e). Vertical lines indicate CTs from the (solid red) NS to S state and (solid green) S to NS state, and almost-occurring CTs from the (dashed red) NS to S state and (dashed green) S to NS state. The model parameters are specified in Table 2. The algorithm parameters are specified in Table 3, the horizontal lines indicate the corresponding thresholds of (red) α and (green) β.

Figure S-1 (a) shows a clear-cut example of a CT between the NS and S states detected by the algorithm. Figure S-1 (b) illustrates a common type of almost-occurring CT where (b1) and (b2) are satisfied but (b3) is not. Another almost-occurring CT is shown in Fig. S-1 (c), however, even though x only exceeds α once after briefly falling below β at the dashed vertical green line, the algorithm still considers the model to be in the S state until (b1)-(b3) are satisfied. Figure S-1 (d) shows that the algorithm does not treat the opposite scenario to (b) in the same way. Regardless of how many $\tilde{t}_1$s are detected, the algorithm still considers the model to be in the NS state until (a1)-(a3) are satisfied. The example in Fig. S-1 (e) illustrates the most common almost-occurring CT from the NS to S state; (a1) and (a2) are satisfied but (a3) is not.

BCTs and BNCTs: We integrate the model from $t = 0$ to $t = 60$ with the parameters specified in Table 2, we set $s = 1$ ($s = -1$)

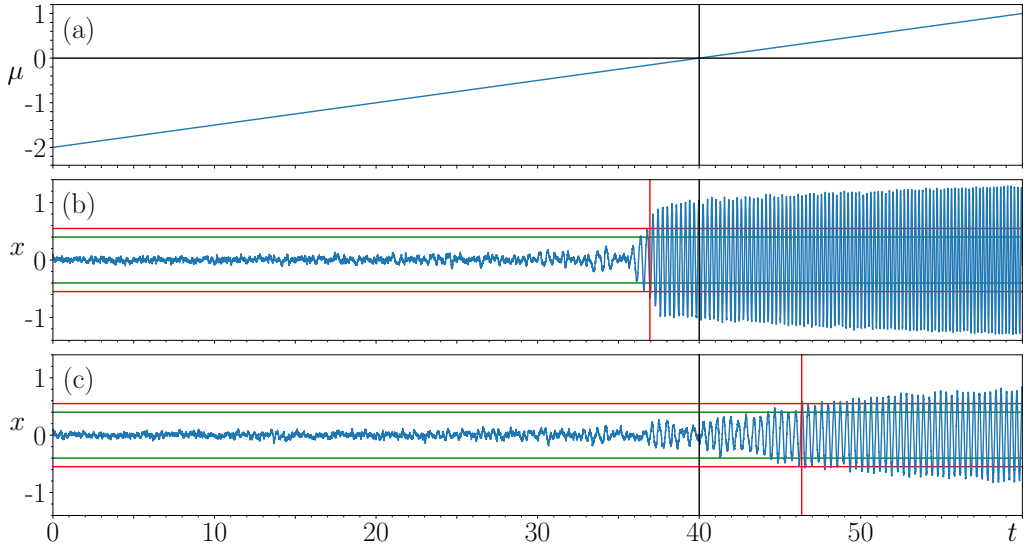

Figure S-2. Examples of BNCTs and BCTs from the NS to S state in the model's output (the variable x from System (5)) that were detected by the algorithm. The evolution of μ versus t is shown in (a), an example of a BNCT is shown in (b) and a BCT is shown in (c). Black vertical lines at $t = 40$ indicate when the respective bifurcations occur. Red vertical lines indicate CTs from the NS to S state (t_1). The model parameters are specified in Table 2. The algorithm parameters are specified in Table 3, the horizontal lines indicate the thresholds of (in red) α and (in green) β.

so that a BNCT (BCT) from the NS to S state can occur. Fig. S-2 (a) shows that μ increases linearly from -2 to 1 according to Eq. (4). We apply our algorithm to (x) the model's output using the parameters values specified in Table 3. In Figs. S-2 (b) and (c) we illustrate an example of a BNCT and BCT detected by the algorithm. The vertical lines indicate (in black) when the bifurcation takes place and (in red) when the algorithm detects a CT from the NS to S state, i.e., the value of t_1. Figure S-2 (b) shows the algorithm may detect a BNCT before the bifurcation takes place. The opposite is shown in Fig. S-2 (c) for the BCT. We study this difference between values of t_1 in greater detail in Fig. S-7 in S4.

S2: Applying the algorithm to detect real seizure activity in the voltage recordings

In this subsection we illustrate that our algorithm detects CTs between the NS and S states in the voltage recordings at times which closely correspond with the expert's annotations. We also describe adjusting our algorithm to account for artefacts that sometimes appear in voltage recordings but are excluded from expert annotations.

Algorithm vs. expert: We first apply our algorithm to the portion of voltage recordings shown in Fig. 4 (b) in the Results section. To provide a comparison to the CTs detected in Figs. S-1 and S-2, we use the same τ_w, Δ, τ_S, and τ_{NS}, and set $\alpha = 0.055$ and $\beta = 0.04$ since both the NS and S states in the model's output are approximately an order of magnitude larger than the corresponding states in the voltage recordings (see values on vertical axes). In Fig. S-3 we compare the values of t_1 and t_2 obtained by the algorithm to the corresponding times provided by the expert. The vertical lines indicate the values of t_1 and t_2 according to (in orange and purple) the expert and (in red and green) the algorithm. The horizontal lines indicate the values of (in red) α and (in green) β used by the algorithm. Figure S-3 shows that, for this choice of algorithm parameters, the expert's t_1 is less than the algorithm's, and the algorithm's t_2 is slightly less than the expert's. In this case, we say there is partial overlap between the seizure intervals defined by the algorithm's and expert's seizure onset and offset times.

Accounting for artefacts: We now describe how we alter our algorithm to account for artefacts that sometimes appear in the voltage recordings; see M1 of the Methods section for reasons why artefacts appear. In Fig. S-4 we show a typical example of artefact activity for $t \in [2, 5]$. We observe that, during this time interval, the LFP jumps between -0.15 and 0.15 mV in much short time intervals, sometimes within consecutive measurements (0.001s). The algorithm may consider this as a CT from the NS to S state since α is crossed and the LFP continues to exceed β for longer than τ_S. Based on this observation, we adapt our algorithm to account for artefacts by altering (a2) as follows:

(a2-1) $|x(t)|$ exceeds α at time $t = t_1$ **and** $|x(t+\delta) - x(t)| < \xi$ for all $t \in [t_1, t_1 + t_w]$, or,

(a2-2) $|x(t)|$ exceeds α at time $t = t_1$ **and** $|x(t+\delta) - x(t)| \geq \xi$ for any $t \in [t_1, t_1 + t_w]$.

If (a2-1) is true then the algorithm continues as before to evaluate whether a CT from the NS to S state occurs at t_1. On the

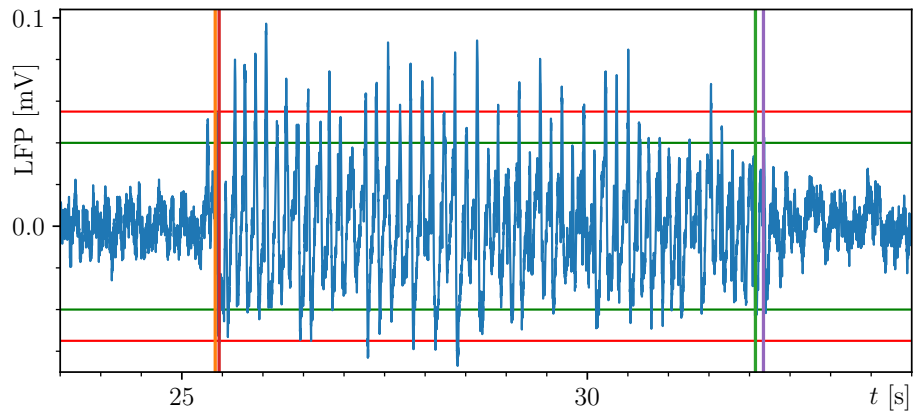

Figure S-3. Expert annotations vs. Algorithm. Vertical lines indicate values of t_1 and t_2 according to the (orange and purple) expert's annotations and (red and green) algorithm for the same portion of the voltage recordings shown in Fig. 4 (b). Algorithm parameters used: $\alpha = 0.055$, $\beta = 0.04$, $\tau_{\mathrm{NS}} = 5$, $\tau_{\mathrm{S}} = 2$, $\tau_w = 1$, and $\Delta = 0.001$, horizontal lines indicate the thresholds of (red) α and (green) β.

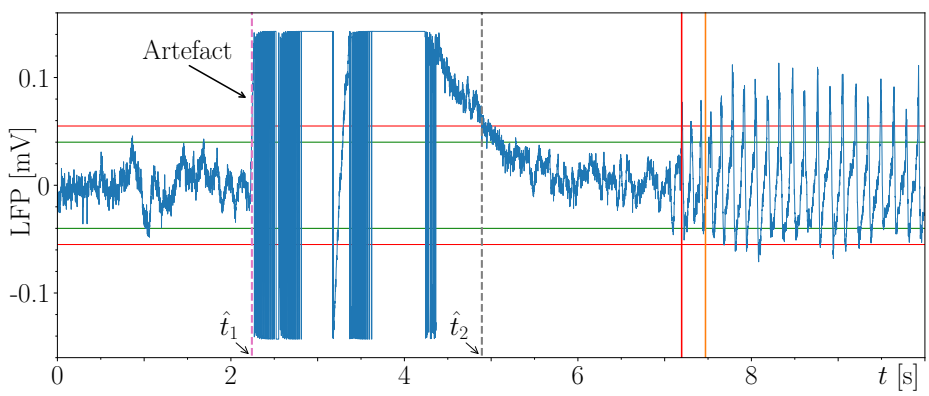

Figure S-4. Accounting for artefacts in the NS state. Vertical dashed lines indicate the (pink) beginning and (grey) end of artefact activity according to our algorithm, i.e., the values of $\hat{t}_1$ and $\hat{t}_2$. Vertical solid lines indicate values of t_1 according to the (orange) expert's annotations and (red) algorithm when applied to a portion of the voltage recordings from session 'T8C'. Algorithm parameters used: $\alpha = 0.055$, $\beta = 0.04$, $\tau_{\mathrm{NS}} = 5$, $\tau_{\mathrm{S}} = 2$, $\tau_w = 1$, and $\Delta = 0.001$, horizontal lines indicate the thresholds of (in red) α and (in green) β.

other hand, if (a2-2) is true, meaning that the difference between two consecutive measurements is greater than $\xi > 0$, then we say artefact activity begins at time $t = \hat{t}_1 = t_1$ and we continue to monitor $|x(t)|$ in moving windows. We say artefact activity ends at time $t = \hat{t}_2$ if $|x(\hat{t}_2)| < \alpha$ **and** $|x(t+\delta) - x(t)| < \xi$ for all $t \in \left[\hat{t}_2, \hat{t}_2 + t_w\right]$. We find most artefacts are accounted for by choosing $\xi = 0.2$.

While artefacts can appear at various locations in the voltage recordings, the example in Fig. S-4 is chosen to show that for the current choice of τ_{S} and τ_{NS}, the algorithm would have considered the time when artefact activity begins as the time when a CT from the NS to S state occurs. Furthermore, this alteration to the algorithm does not diminish the strong agreement between the expert and the algorithm; see the corresponding t_1 values in Fig. S-4.

Artefacts nearby CTs: After artefact activity has ended we find the following two scenarios unfold, $|x(t)|$ either remains below α or quickly re-exceeds α. The second scenario indicates that the artefact activity happens in conjunction with (i) additional artefact activity, (ii) an almost-occurring CT from the NS to S state, or most importantly (iii) a CT from the NS to S state. Fig. S-5 shows that for (iii), minor disagreements between the algorithm and the expert arise in terms of t_1. More specifically, Fig. S-5 shows an example of when the expert's t_1 (vertical orange line) is similar to $\hat{t}_1$ (vertical dashed pink line), the time where the algorithm says that artefact activity begins. It is only after the artefact activity ends (vertical dashed grey line) that the algorithm says there is a CT from the NS to S state (vertical red line).

While one may argue that further alterations should be made to the algorithm to improve the agreement between the algorithm

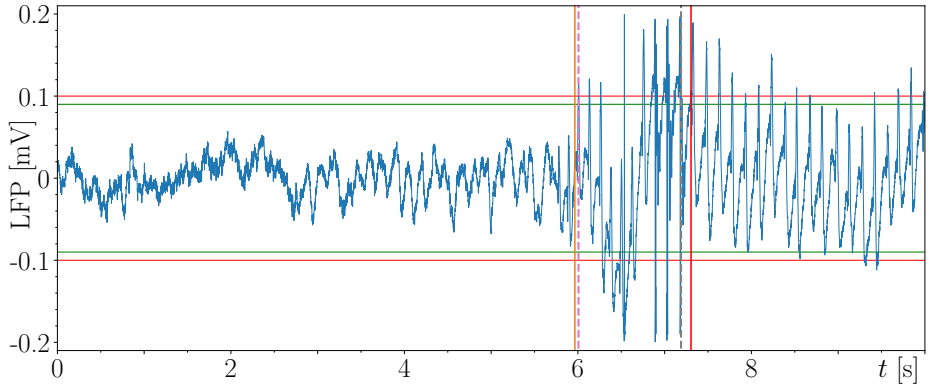

Figure S-5. Accounting for artefacts in the NS state which happen in conjunction with a CT. Vertical dashed lines indicate the (pink) beginning and (grey) end of artefact activity according to our algorithm, i.e., the values of $\hat{t}_1$ and $\hat{t}_2$. Vertical solid lines indicate values of t_1 according to the (orange) expert's annotations and (red) algorithm when applied to a portion of the voltage recordings from session 'K7E'. Algorithm parameters used: $\alpha = 0.1$, $\beta = 0.09$, $\tau_{\mathrm{NS}} = 3$, $\tau_{\mathrm{S}} = 2$, $\tau_w = 1$, and $\Delta = 0.001$, horizontal lines indicate the thresholds of (in red) α and (in green) β.

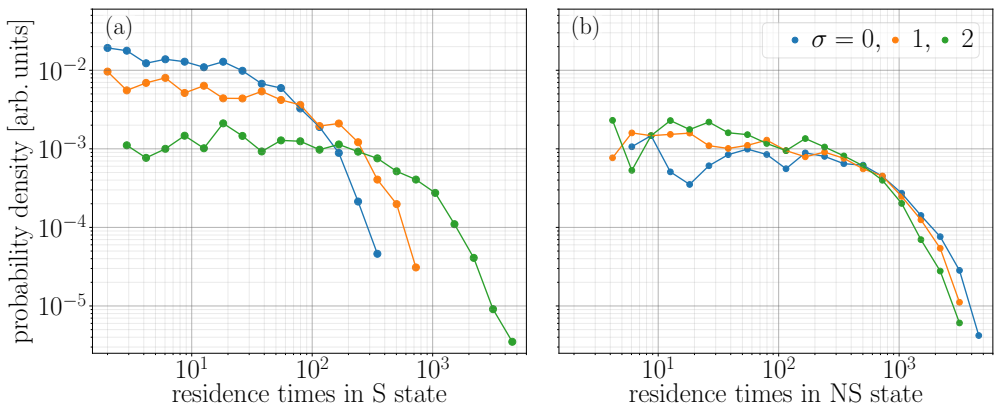

Figure S-6. Probability density of residence times in the model's (a) S state and (b) NS state for the values of σ specified in the plot legend. The information presented here is based on 700 CTs between the NS and S states in the model as detected by the algorithm using the parameter values specified in Tables 2 and 3.

and the expert's annotations in such instances, for the purposes of this paper (classifying seizure generation mechanisms), it is more important that we detect this artefact activity in order to prevent the SVM from classifying the corresponding CT. See S10 for further details.

S3: Influence of shear on residence times in NS and S states for noise-induced CTs

In this subsection we examine how σ, the shear parameter in our model, influences the model's residence times in the NS and S states. More specifically, for $\sigma = 0, 1$, and 2, we integrate the model forward in time up to $t = 1 \times 10^6$ with the remaining model parameters as specified in NCT section of Table 2. We apply our algorithm to x, the model's output, using the algorithm parameter values specified in Table 3. Our algorithm detected the following number of CTs between the NS and S states for different values of σ: 948 CTs for $\sigma = 0$, 984 CTs for $\sigma = 1$, and 715 CTs for $\sigma = 2$. We then compute residence times in the NS and S states based on the t_1 and t_2 values obtained for each choice of σ. Since the number of CTs vary depending on σ, and we would like to compare residence times across different values of σ, we compute the probability density of residence times in the NS and S states for each σ for 700 randomly chosen residence times, reflecting the smallest number of CTs detected when $\sigma = 2$. The resulting probability densities of residence times in the S state are shown in Fig. S-6 (a) and NS state in Fig. S-6 (b).

Figures S-6 (a) and (b) show that the model is most likely to spend a short duration of time in a given state and least likely to spend a large duration of time. Figure S-6 (a) shows that σ has a strong influence on residence times in the S state, and Fig. S-6 (b) shows it has no influence on the NS state. More specifically, Fig. S-6 (a) shows that as σ increases the probability that the model spends (i) a relatively short duration of time in the S state slightly decreases, and (ii) a relatively large duration

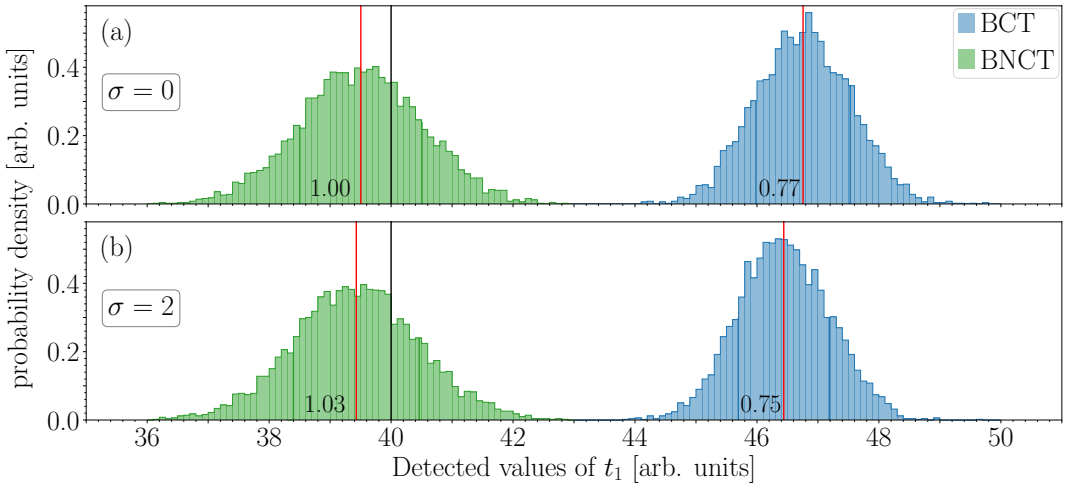

Figure S-7. Probability density of t_1 values for (blue) BCTs and (green) BNCTs with shear parameter set to (a) $\sigma = 0$ and (b) $\sigma = 2$. The vertical black line indicates when the respective bifurcations take place, vertical red lines indicates the mean t_1 for each CT-type, and the vertical blue and green lines indicate the corresponding standard deviation (quoted the left of each red line). The information presented here is based on 5000 examples of each CT-type.

of time in the S state increases significantly; for the largest recorded residence time in the S state when $\sigma = 0$, there is an approximately equal probability that the model is in the S state for nearly ten times longer when $\sigma = 2$.

S4: Influence of shear on the detected t_1 values for bifurcation and bifurcation/noise induced CTs

In this subsection we examine how σ, the shear parameter in our model, influences when BCTs and BNCTs from the NS to S state occur. More specifically, for $\sigma = 0$ and 2, we generate 5000 different examples of BCTs and BNCTs for the model parameters specified in Table 2. We apply our algorithm to x, the model's output, using the algorithm parameter values specified in Table 3. We then construct probability density diagrams (based on histograms) of the different t_1 values detected by our algorithm, these are shown in Fig. S-7 (a) for $\sigma = 0$ and Fig. S-7 (b) for $\sigma = 2$ where the data in blue corresponds to BCTs and in green corresponds to BNCTs. In both Figs. S-7 (a) and (b) the vertical black line indicates when both bifurcations happen, the vertical red lines indicate the mean t_1 for both BCTs and BNCTs, the vertical blue and green lines indicate the corresponding standard deviation (which is quoted to the left of each red line).

BNCTs: Figure S-7 shows that as σ increases there are relatively small changes in the statistics of the detected t_1 values; the mean shifts slightly to the left and the standard deviation slightly increases. For both σ values the mean t_1 value is relatively close when the bifurcation happens. From additional experiments we found that decreasing ν results in the mean t_1 value increasing. If ν is small enough then the CTs becomes purely bifurcation-induced. If ν is too large then few CTs will take place after the bifurcation. The chosen value of ν in Table 2 enables the model to produce BNCTs with an almost equal likelihood that noise or the bifurcation triggering the CT.

BCTs: Figure S-7 shows that as σ increases there are relatively small, but more noticeable, changes in the statistics of the detected t_1 values than BNCTs. Specifically, the mean t_1 value and standard deviation both decrease.

Two aspects of Fig. S-7 are critical for the main aim of the paper of classifying CTs. The first is that there is no overlap between the different distributions; BNCTs are designed to cluster around the time of bifurcation while BCTs are designed to occur at a much greater time after the bifurcation. The second, there is no $t_1 < 35$; this informs our choice of $T^- = -30$ when producing Fig. 6, allowing us to experiment with the relatively wide range of t_m values in the interest of finding the optimal way to differentiate between the three seizure generation mechanisms via our SVM classifier.

S5: Tuning algorithm parameters

In this subsection we describe the procedure used to tune the algorithm's parameters to maximise agreement between the algorithm and the expert's annotation. By agreement we mean that, in the sense of a receiver-operator-characteristic (ROC) analysis, the proportion of correctly labelled seizure intervals is maximal and the proportion of incorrectly labelled non-seizure intervals is minimal, using the expert's annotations as the reference. Seizure and non-seizure intervals are defined by seizure onset and offset times provided by the expert and the values of t_1 and t_2 obtained by the algorithm when applied to a given set of voltage recordings. We consider a seizure interval defined by the algorithm's t_1 and t_2 values to be correctly labelled if it (i) is contained within, (ii) contains, or (iii) partially overlaps with a seizure interval defined by the expert's annotations. See

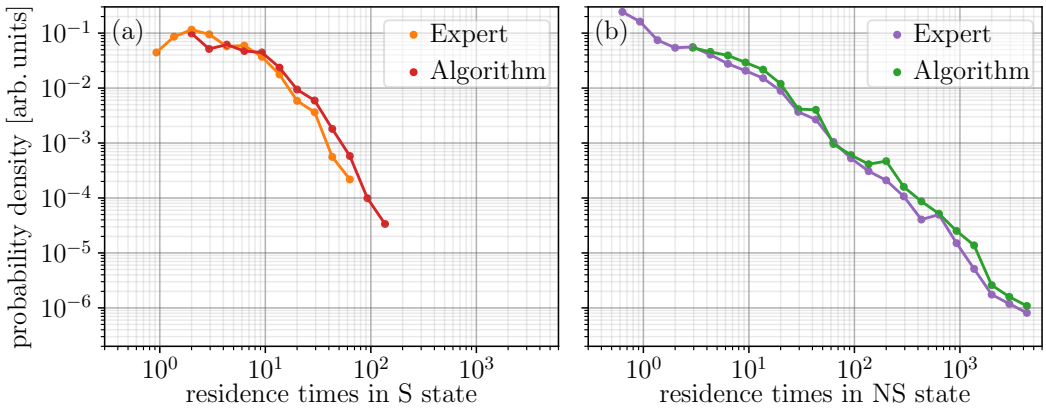

Figure S-8. Probability density of residence times in the (a) S state and (b) NS state for rat S according to the (orange and purple) expert's annotations and (red and green) the algorithm.

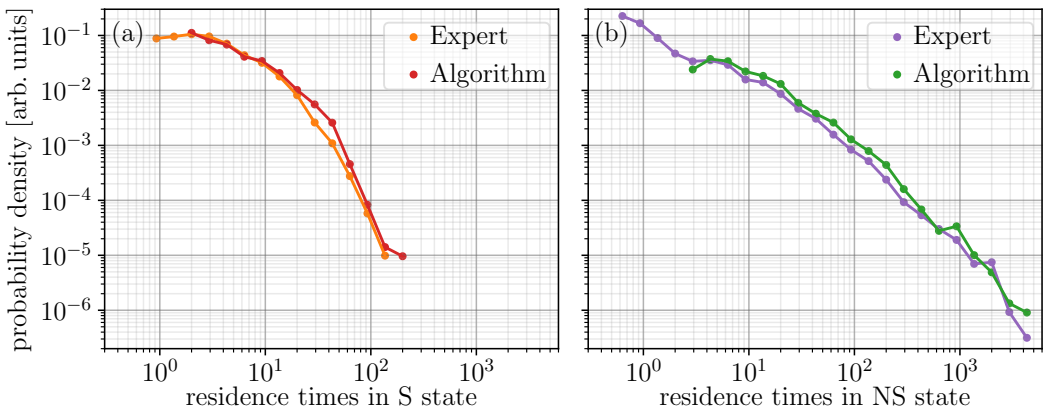

Figure S-9. Probability density of residence times in the (a) S state and (b) NS state for rat T according to the (orange and purple) expert's annotations and (red and green) the algorithm.

Fig. S-3 in S2 for an example of (iii). Similarly, we consider an incorrectly labelled non-seizure interval to be a seizure interval defined by the algorithm's t_1 and t_2 values that is contained within a non-seizure interval defined by the expert's annotations.

From several experiments we found that the above agreement is most sensitive to different values of α. Therefore, for each recording session, we select a different α from $[0.03, 0.1]$ to obtain the best agreement with the expert annotations. β is adjusted accordingly so that $\beta = \alpha - 0.01$. The values of τ_w, Δ, τ_S, and τ_{NS} specified in Table 3 provide a strong agreement between the algorithm and expert across all the analysed voltage recordings from rats S, T, and K. This strong agreement is illustrated in Figs. S-8, S-9, and S-10 in terms of the probability distributions of residence times in the NS and S states for rats S, T, and K.

Figure S-8 shows there is relatively little difference between the probability distributions of residence times in the NS and S states according to the expert and the algorithm when computed from all the voltage recordings of rat S. Small differences occur at the end points between the corresponding distributions, this is mainly due to the choice of τ_{NS} and τ_S being less than the smallest residence times in the NS and S states according to the expert's annotations. As a result, a given seizure/non-seizure interval according to the algorithm may consist of multiple seizure and non-seizure intervals according to the expert. However, when τ_{NS} and τ_S are more in line with the expert's annotations, we find this significantly reduces the agreement between the algorithm and the expert. Figures S-9 and S-10 show that similar results are obtained for rats T and K.

S6: Additional clarification on how TSPs and m(TSPs, t_m) are computed for CTs in the model's output

We obtain 100 BCTs, 100 BNCTs, and 100 NCTs from the model's NS to the S state as follows. For NCTs, we use the model parameters specified in Table 2, algorithm parameters specified in Table 3, and integrate the model until our algorithm detects 100 CTs between the NS and S states that satisfy the following criteria: the model is in the NS state for all $t \in [t_1 + T^-, t_1)$ and is in the S state for all $t \in [t_1, t_1 + T^+]$ where $T^- = -30$ and $T^+ = 10$. For BCTs and BNCTs, we use the model parameters

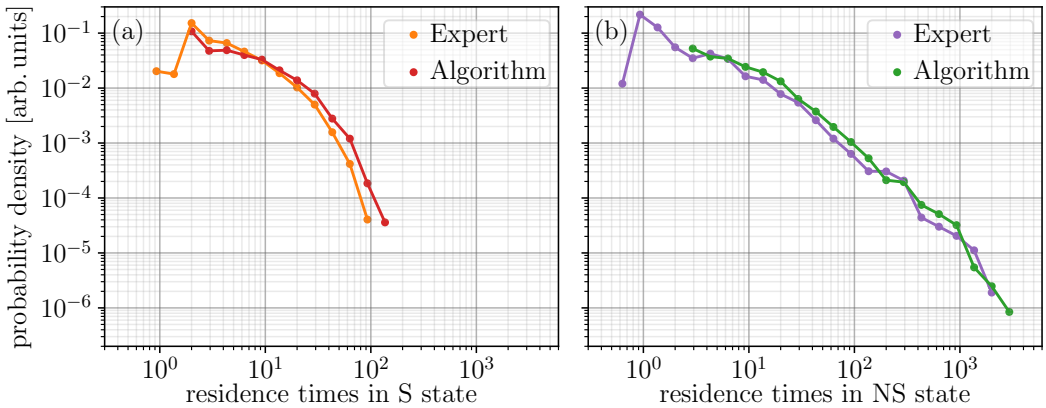

Figure S-10. Probability density of residence times in the (a) S state and (b) NS state for rat K according to the (orange and purple) expert's annotations and (red and green) the algorithm.

specified in Table 2, different noise realisations for each simulation of the model, and detect CTs using algorithm parameters specified in Table 3. In Fig. S-7 in S4 we show that all these BCTs and BNCTs satisfy the above two conditions as each $t_1 > 30$ and $t_2 - t_1 > 10$. Further, we only consider CTs where no almost-occuring CTs appear for $t \in \left[t_1 + T^-, t_1\right)$.

Based on the above we obtain 100 examples of BCTs, NCTs, and BNCTs from the model. When computing the TSPs and m(TSPs, t_m) near each CT, it is convenient to work with a new time $T = t - T - 1$ where $T \in \left[T^-, T^+\right] = [-30, 10]$ and each CT is detected at $T = 0$. We then follow the steps in M2 of the Methods section to compute the TSPs and m(TSPs, t_m) for a given $t_w = 1$ and chosen values of t_m. This choice of T^- and T^+ results in the GV, $\log_{10}$GV, and AC being defined from $T = T^- + t_w = -29$, and $\log_{10}$GV(AC) is defined from $T = T^- + 2t_w = -28$. Further, the first values of m(GV, t_m), m($\log_{10}$GV, t_m), and m(AC, t_m) are defined from $T = T^- + t_w + t_m$ and m($\log_{10}$GV(AC), t_m) is defined from $T = T^- + 2t_w + t_m$. Thus, if $t_m = 8$ then m(GV, 8) is defined from $T = -21$ and m($\log_{10}$GV(AC), 8) is defined from $T = -20$, and so on.

S7: TSPs for a single example of each CT type

In this subsection we show how the time series properties (TSPs) of the model's output behave nearby an example of a BCT, BNCT, and NCT from the NS to S state. Each of these TSPs are specified in Table 5. We generate and detect the BCT and BNCT in Figs. S-11 (a) and (b) using the same parameters as in Fig. S-2. For the NCT case in Fig. S-11 (c) we generate and detect a CT for the same parameters from Fig. S-1 and consider a similar example of an NCT to Fig. S-1 (a). For each of the CTs shown in Figs. S-11 (a)-(c), we plot the model's output, x, versus $T = t - t_1$ for the corresponding values of t_1 detected by our algorithm. We analyse the TSPs of each CT from $T = -30$ in small moving windows of length $t_w = 1$. Thus, for the reasons outlined in M2 of the Methods section, the GV, $\log_{10}$GV, and AC are defined from $T = -29$, and $\log_{10}$GV(AC) is defined from $T = -28$. We now discuss how each of these TSPs behave in relation to the different CT types.

GV: Figures. S-1 (d) to (f) show that the GV of the model's output in the S state is much larger than the NS state for each CT-type. One would naturally expect this to happen when a transition from a small to a large-amplitude of oscillation occurs. In the BCT case (Fig. S-11 (d)), there is a noticeably large and relatively constant increase in GV from $T \approx -7$ onwards, which is just before the bifurcation point indicated by the inward black tick. For the BNCT and NCT cases (Figs. S-11 (e) and (f)), no significant change in GV is seen before the CTs are detected. However, after the CTs are detected, the increase in GV is much greater for the BNCT case than the NCT case. This difference arises because the radius of the limit cycle for the BNCT case is larger than for the NCT case (see Eq. (3)), thereby resulting in a more dramatic change in the model's output.

$\log_{10}$GV: Figures S-11 (g) and (h) show $\log_{10}$GV increases at a relatively small but constant rate before the CT in the BCT and BNCT cases, whereas Fig. S-11 (i) shows no such increase for the NCT case. Similar to the GV results, there is a sudden and large increase after the CT for both the BNCT and NCT cases. Notably, for the BCT case, $\log_{10}$GV starts to increase at an earlier time than GV in Fig. S-11 (a), thus demonstrating the benefit of monitoring $\log_{10}$GV in addition to GV.

AC: Figures S-11 (j) and (k) show AC increases at a relatively constant rate before the BCTs and BNCTs. For the BCT case, AC remains close to 1 from $T \approx -7$ (mirroring the trend seen in Fig. S-11 (d)). On the other hand, it is only after the CT that AC remains near 1 for the BNCT case. In stark contrast, Fig. S-11 (l) shows that the AC remains relatively close to 1 for the NCT case, with no trends present long before the CT. Interestingly, there is a local minimum in AC present before each CT. This occurs because, regardless of the AC's value before the CT, the model's output becomes highly correlated with itself after the CT due to the inherent periodicity of the limit cycle, thus causing the AC to quickly increase to 1. We also see less fluctuations in the AC as the CT is approached for the BCT and BNCT cases.

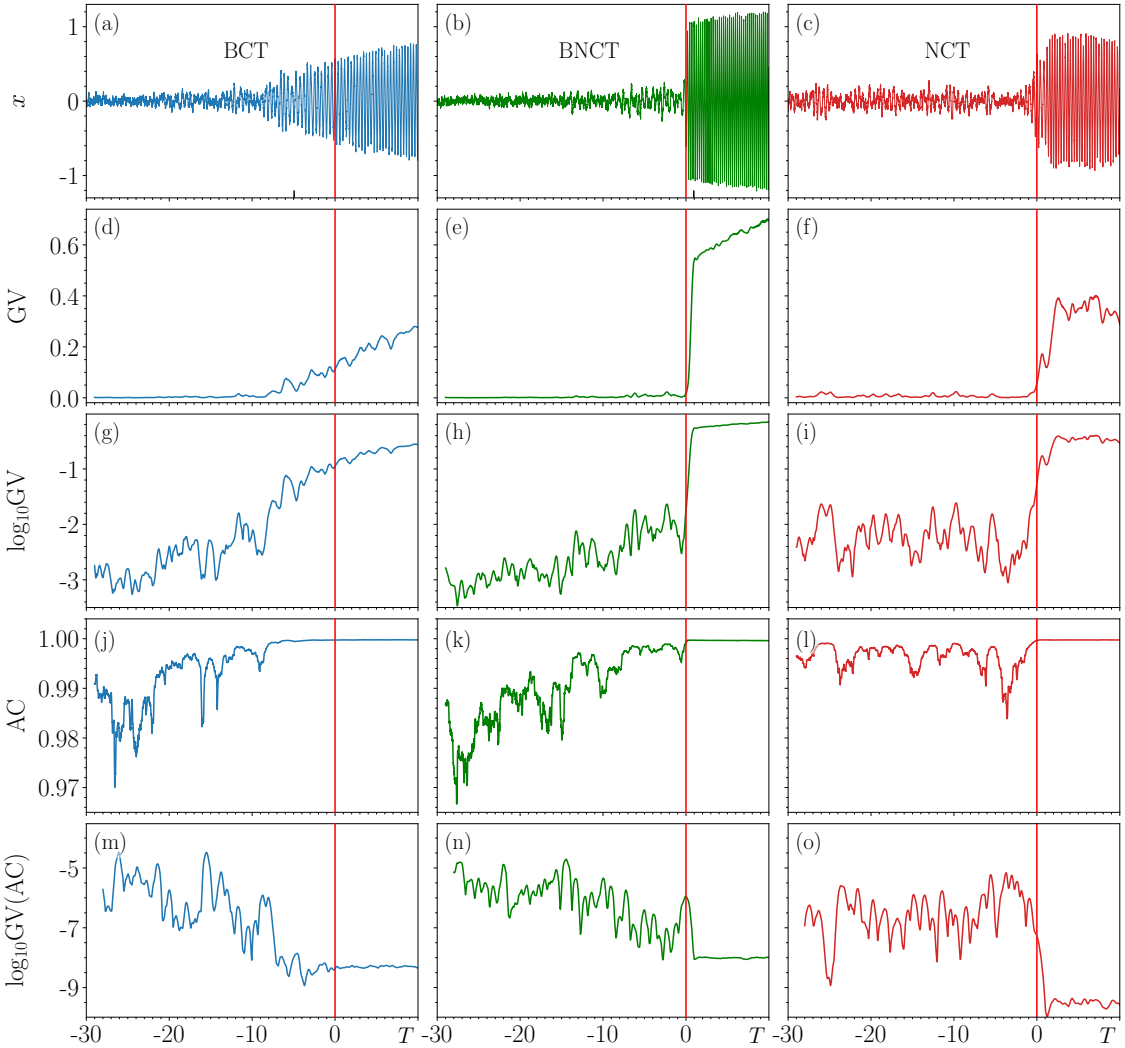

Figure S-11. Time series properties (TSPs) of the model's output near an example of a (left-hand column) BCT, (middle column) BNCT, and (right-hand column) NCT. For each CT-type, the top row shows the model's output versus $T = t - t_1$ with t_1 detected by the algorithm. Inward black ticks in (a) and (b) indicate when the bifurcation occurs. The remaining rows show how each of the TSPs, specified in the vertical axes labels and in Table 5, behave for the same values of T. Vertical red line in each panel emphasises when $T = 0$.

$\log_{10}$GV(AC): Figures. S-11 (m)-(o) provide a more informative picture on the fluctuations mentioned above by examining how $\log_{10}$GV(AC) changes over time. For the BCT and BNCT cases, there is an overall constant decrease in $\log_{10}$GV(AC) before the CT. Log_{10}GV(AC) remains at a similar and constant value after both CTs. For the NCT case, $\log_{10}$GV(AC) remains relatively constant before the CT and quickly decreases to a smaller value after the CT than the BCT and BNCT.

S8: Model's TSPs computed for $t_w = 1$ and 2

In this subsection we provide clarification on comments made earlier regarding tuning the model's parameters so that its output mimics the three important characteristics of real seizure activity in the voltage recordings. In particular, we mentioned that we compromise on characteristic (iii), "Intrinsic time scales of the NS and S states.". We now show that there are little consequences from this in terms of the main aim of this paper, to identify the CT-type responsible for generating a given seizure in voltage recordings.

In Fig.4 (c) and (d), we showed that the timescale of oscillation of the model's S state is approximately half that of the dominant timescale of real seizure activity. Thus, when choosing a t_w to compute the TSPs (the time duration of the moving window), the value of t_w for the model should be twice that of the value used for the voltage recordings. However, we show in Fig. S-12 that there are no major differences in the TSPs of the model when computed using either $t_w = 1$ (in blue) or $t_w = 2$ (in

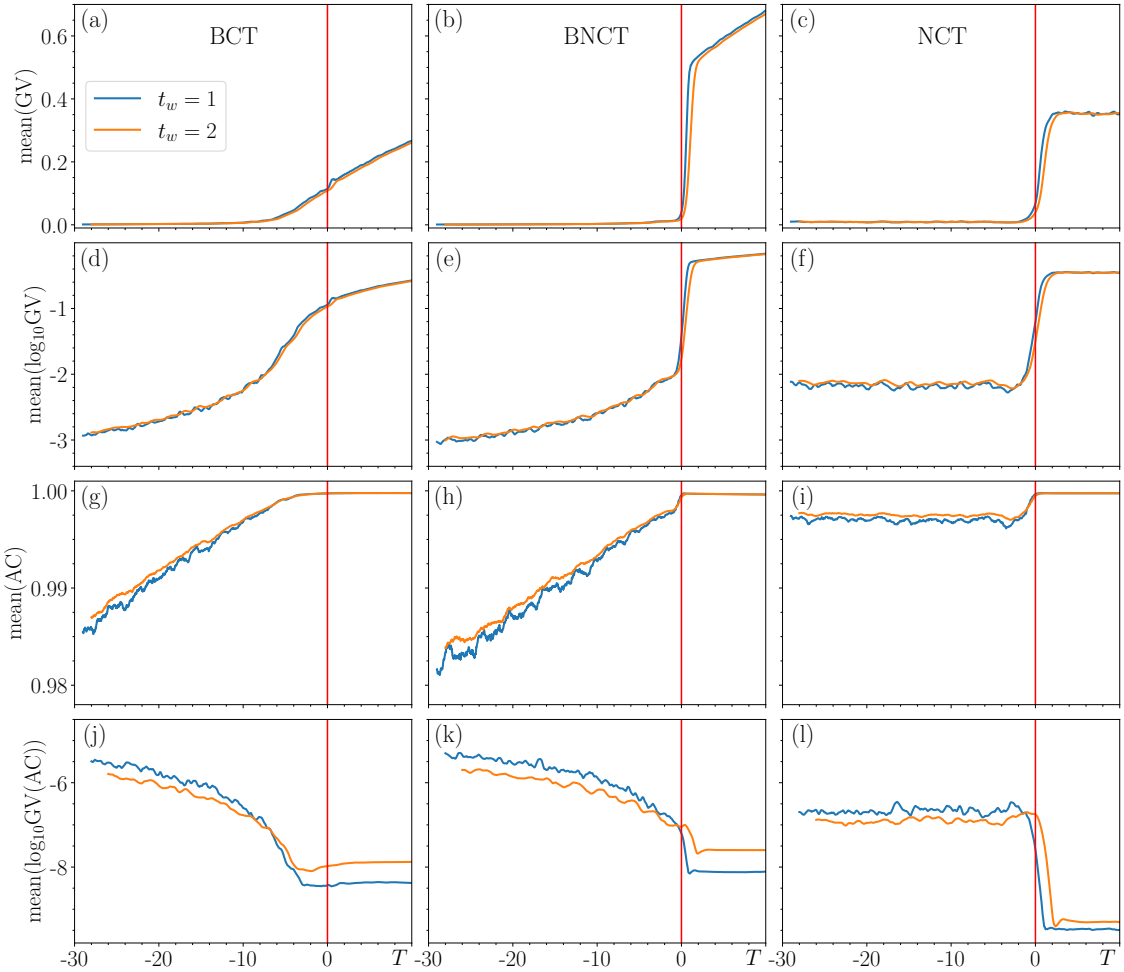

Figure S-12. Comparing different choices of t_w when computing the time series properties (TSPs) of model's output near BCTs (left-hand column), BNCTs (middle column), and NCTs (right-hand column). Each panel shows the mean of each TSP versus $T = t - t_1$ for the t_1 values detected by the algorithm. Each curve represents the mean computed from 100 examples of each CT-type for a given T.

orange). We also find that whether the SVM is trained using TSPs that are calculated with $t_w = 1$ or 2, there is little impact on classifying seizure generation mechanisms in the voltage recordings. For convenience, we choose $t_w = 1$ when computing TSPs from both the model and the voltage recordings.

S9: Computing SVM feature importance via the mean permutation importance (MPI)

In this subsection we follow the steps outlined in M3 of the Methods section on how to compute the 'mean permutation importance' (MPI) in order to quantify the importance of each feature used to construct SVM type-1, 2, and 3 (see Table 6 for details on which features are used to construct each SVM). The results discussed in this subsection serve to accompany the results shown in Fig. 8 where SVM accuracy versus T is shown for SVM type-1, 2, and 3 with $t_m = 4, 8$, and 12. Since we trained and tested a new SVM for different values of T in Fig. 8, we compute the MPI the corresponding values of T.

Figures S-13 and S-14 show the resulting MPI versus T. Each of the solid and dashed curves are coloured to correspond with the features specified in the legend. These curves are surrounded by a cloud of the corresponding colour which describes the standard deviation of the MPI for a given feature. Note, the results shown in Fig. S-13 (c) informed the choice of features that were used to construct Fig. 7.

Figures S-13 and S-14 show that for each SVM, (i) some features are much more important than others and (ii) the importance of each feature varies depending on the value of T, (iii) and also varies depending on the value of t_m in the case of SVM type-2 and 3. For instance, when comparing the results shown in Figs. S-13 (a) and (e) for $T > 5$, Fig. S-13 (a) shows that when $t_m = 4$ the most important feature is $m(\log_{10}\text{GV},\ t_m)$ and least important feature is $m(\log_{10}\text{GV(AC)},\ t_m)$, whereas Fig. S-13 (c) shows

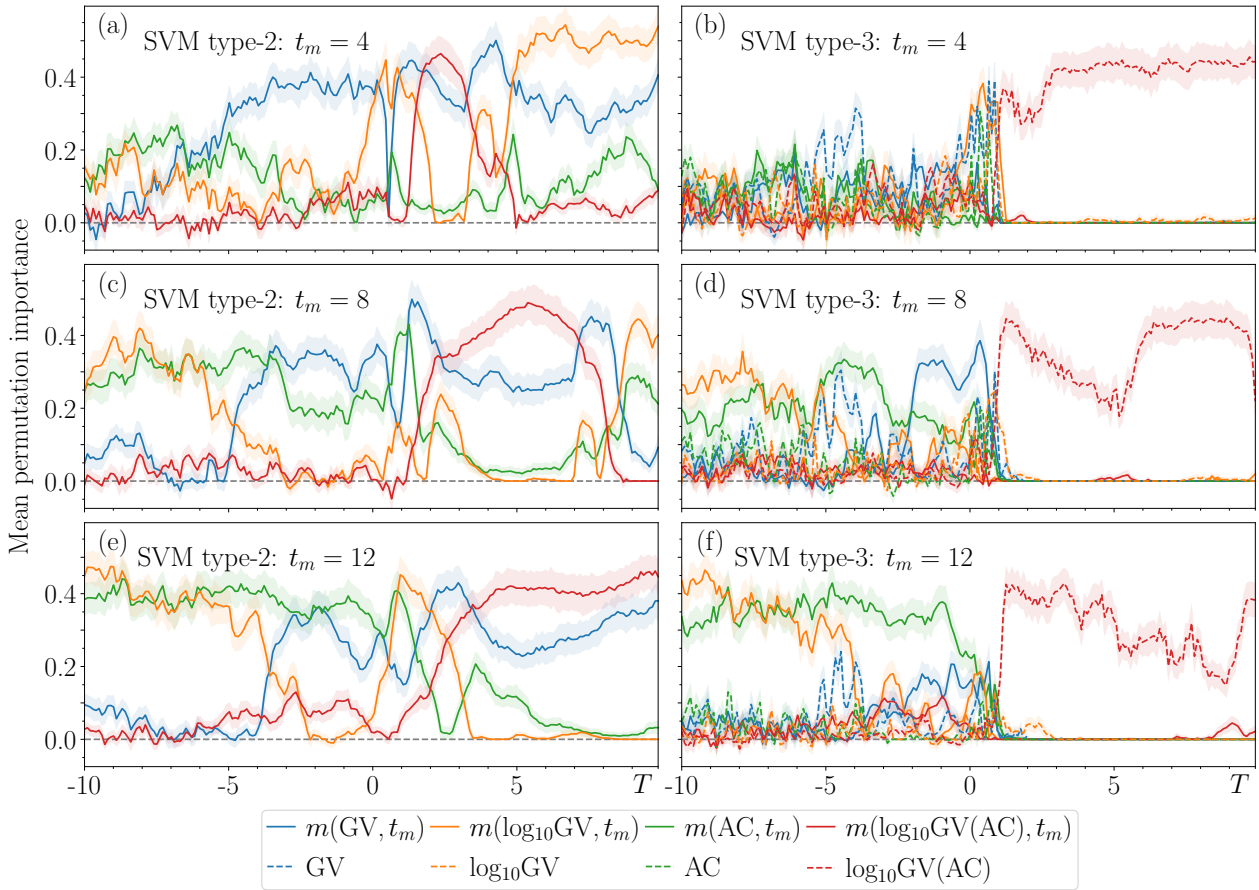

Figure S-13. Feature importance for SVMs type-2 and 3: Mean permutation importance (solid and dashed curves) and its standard deviation (cloud surrounding each curve) at time $T = t - t_1$ for a given feature specified in the plot legend for (left-hand column) SVM type-2 and (right-hand column) SVM type-3 with (a) and (b) $t_m = 4$, (c) and (d) $t_m = 8$, and (e) and (f) $t_m = 12$.

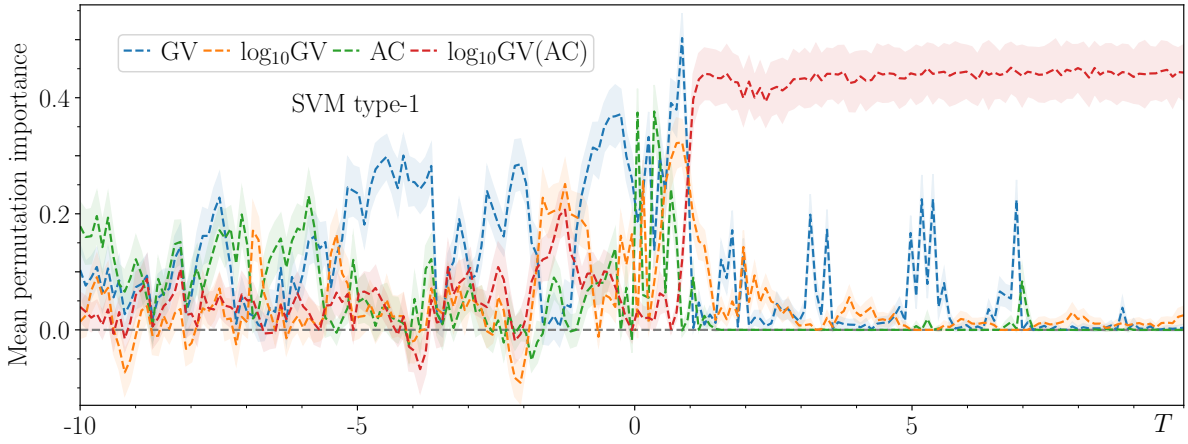

Figure S-14. Feature importance for SVM type-1: Mean permutation importance (dashed curves) and its standard deviation (cloud surrounding each curve) at time $T = t - t_1$ for a given feature specified in the plot legend.

the opposite for $t_m = 12$.

It is also worth noting the significant changes in the MPI of some features near $T = 0$ for SVM type-2 and 3. For example, Fig. S-13 (a) shows that $m(\mathrm{GV},\ t_m)$ suddenly changes from being the most important to least important and back to the most

important feature over a small range of T values near 0. On the other hand, Fig. S-13 (a) shows that $m(\log_{10}\text{GV}, t_m)$ and $m(\log_{10}\text{GV(AC)}, t_m)$ change from being the least important to most important and back to least important feature over a small range of T values between 0 and 5. Similar behaviour is seen in Figs. S-13 (c) and (e) but to a much less significant extent.

While Figs. S-13 (a), (c), and (e) show that the most important feature for SVM type-2 often varies for $T > 0$, a different picture emerges for SVM type-3 in Figs. S-13 (b), (d), and (f) as $\log_{10}\text{GV(AC)}$ is the most importance feature for each choice of t_m when $T > 0$. Similar behaviour is shown in Fig. S-14 for SVM type-1. Figures S-13 (b), (d), and (f) also show that the $m(\text{TSPs}, t_m)$ play a more significant role in the accuracy of SVM type-3 for larger t_m before the CT. This further illustrates the effects of combining the TSPs and their slopes in SVM type-3. More specifically, for $T < 0$ the features which are most important for SVM type-2 are the most important for SVM type-3, for $T > 0$ the features which are most important in SVM type-1 are the most important for SVM type-3.

S10: TSPs and m(TSPs, t_m) for CTs not suitable for classification

In this subsection we show why it is not suitable to classify CTs from the NS to S state when artefacts or almost-occurring CTs happen in close proximity to the CT. More specifically, we show that some of the corresponding TSPs and $m(\text{TSPs}, t_m)$ listed in Table 5 deviate significantly from those of a typical CT from the NS to S state when artefacts or almost-occurring CTs happen nearby. This is important to show because if the SVM classifies CTs using these TSPs and $m(\text{TSPs}, t_m)$ then the classifications are much less reliable and may skew the main results presented in Table 4.

In Fig. S-15 we show the examples of CTs that we compute the TSPs and $m(\text{TSPs}, t_m)$ of in this subsection. In each example we plot the LFP versus $T = t - t_1$, where t_1 corresponds to the time where the algorithm detects a CT from the NS to S state. The example in Fig. S-15 (a) (same as in Fig. S-3) is a typical example of a CT that satisfies C1-C5 and is chosen as a reference point to compare with the following examples of CTs that satisfy C1-C3 but dot not satisfy C4 or C5: Fig. S-15 (b) (taken from recording session 'T8C') shows an example of an almost-occurring CT before a CT from the NS to S state, Fig. S-15 (c) (same as in Fig. S-4) shows an example of when artefact activity occurs before a CT from the NS to S state, and Fig. S-15 (d) (same as in Fig. S-5) shows an example of when artefact activity occurs in conjunction with a CT from the NS to S state. Note also that the example in Fig. S-15 (b) provides similar insight to a CT that does not satisfy C1 and a CT generated by the model that does not satisfy the corresponding selection criteria. The horizontal red and green lines in each panel of Fig. S-15 indicate the values of α and β used to detect CTs between the NS and S states, these values are chosen according to the procedure outlined in S5. The vertical red solid lines indicates a CT from the NS to S state. Vertical dashed lines indicate (in red) an almost-occurring CT from the NS to S state, (in pink) the beginning and (in grey) end of artefact activity.

In Fig. S-16 we show how each of the TSPs and $m(\text{TSPs}, t_m)$ with $t_m = 8$ behave versus T for the voltage recordings shown in Fig. S-15. More specifically, the different coloured curves in each panel of Fig. S-16 correspond to a given TSPs or $m(\text{TSPs}, t_m)$ of the examples shown in Fig. S-15;:the blue curves correspond to Fig. S-15 (a), the green curves correspond to Fig. S-15 (b), the orange curves correspond to Fig. S-15 (c), and the red curves correspond to Fig. S-15 (d). The different coloured vertical lines have the same meaning as in Fig. S-15. The pink and grey vertical dashed lines at $T \approx -5$ and -2.25 correspond with the example shown in Fig. S-15 (c), the other pink and grey vertical dashed lines at $T \approx -1.25$ and -0.1 correspond with the example shown in Fig. S-15 (d).

Figure S-16 shows that when $T = 0$ there is an almost immediate change in the TSPs and a slightly delayed and longer lasting change in the $m(\text{TSPs}, t_m)$ in response to the presence of CTs that almost occur and the different artefacts. This delayed and longer lasting change in the $m(\text{TSPs}, t_m)$ results in major differences from the $m(\text{TSPs}, t_m)$ of a typical CT without any artefacts or almost-occurring CTs nearby, i.e., when comparing the blue curves to the others. These differences are visible at several values of T and, crucially for our SVM setup, are noticeable for $T = 2$ which is the value of T we base our final results on in Table 4, i.e., the value of T where our SVM classifies CTs in the voltage recordings. In the following headings we highlight these differences.

Almost-occurring CT: (Comparing blue and green curves at $T = 2$) Considerable differences particularly in $m(\log_{10}\text{GV}, t_m)$ and $m(\text{AC}, t_m)$.

Artefact before CT: (Comparing blue and orange curves at $T = 2$) Major differences particularly in $m(\text{GV}, t_m)$ and $m(\log_{10}\text{GV}, t_m)$ but quite similar in $m(\text{AC}, t_m)$ and $m(\log_{10}\text{GV(AC)}, t_m)$.

Artefact in conjunction with CT: (Comparing blue and red curves at $T = 2$) Major differences in all $m(\text{TSPs}, t_m)$, there appears as if no changes happens in $m(\text{AC}, t_m)$ and $m(\log_{10}\text{GV(AC)}, t_m)$.

As a final word, we would like to add the following for clarification. Since the differences highlighted above are only present in the $m(\text{TSPs}, t_m)$ when $T = 2$ one may argue that in the interest of classifying as many CTs in the voltage recordings as possible, one should do so using SVM type-1 instead of SVM type-3 so that less CTs like those shown in Fig. S-15 need to be excluded via conditions C1-C5. However, we found that SVM type-1 provides classifications which are in total disagreement with the fundamental characteristics of the different CTs that are captured by the $m(\text{TSPs}, t_m)$; there were instances when some CTs were classified as BCTs without the typical the increase in GV and AC before the CT, and other CTs were classified as

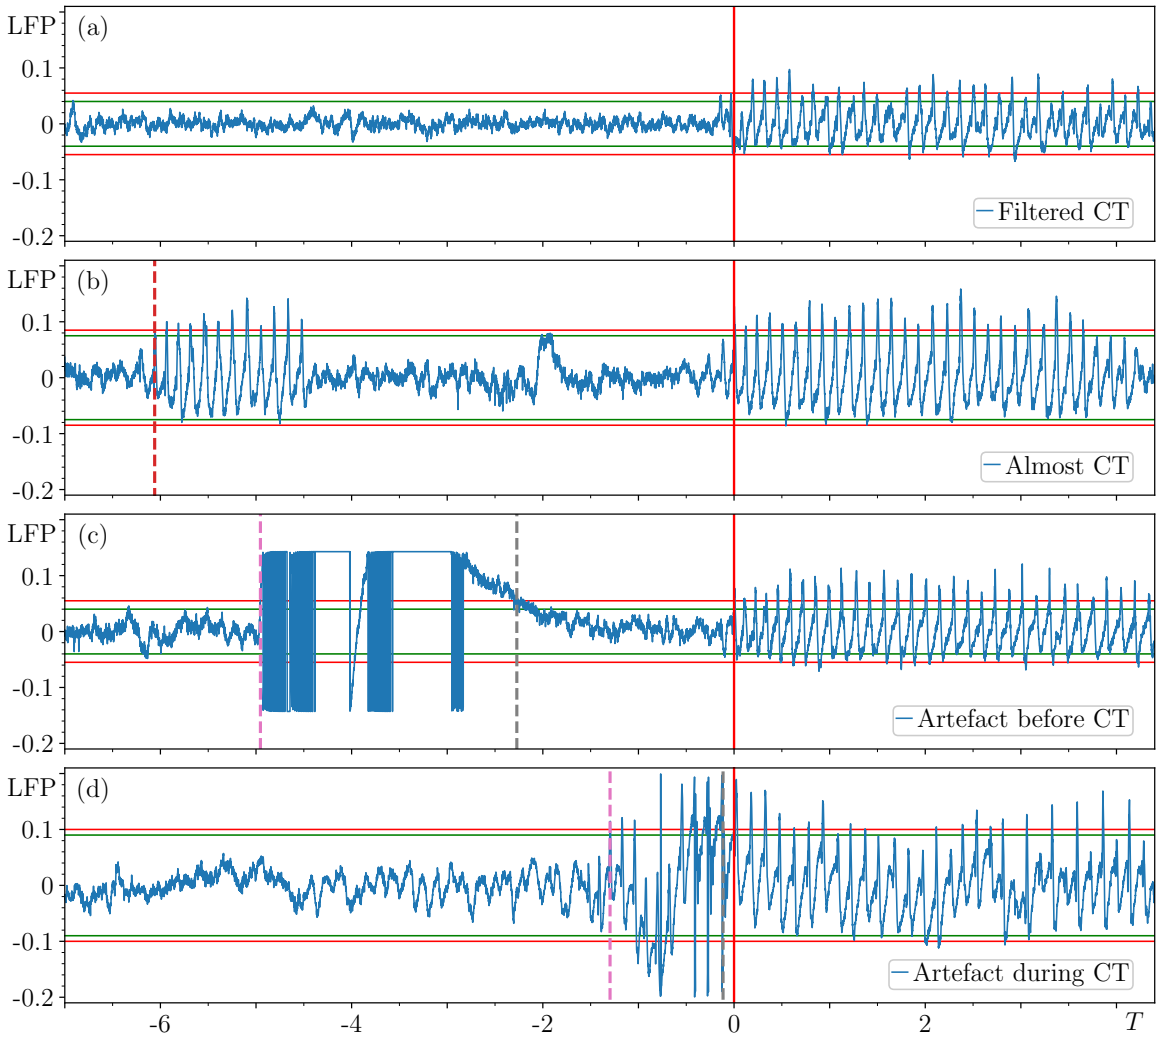

Figure S-15. Examples of detected CTs in voltage recordings that (a) are suitable for classification, (b)-(d) are not suitable for reasons given in each plot legend, and are plotted versus $T = t - t_1$ for t_1 values detected by our algorithm. Note, in (d), the 'artefact during CT' label refers to the 'artefact in conjunction with a CT' scenario. The vertical lines indicate (solid red) when $T = 0$, (dashed red) an almost-occurring CT from the NS to S state, (dashed pink and grey) the beginning and end of artefact activity, according to our algorithm. Horizontal solid lines indicate thresholds of (red) α and (green) β used by the algorithm.

NCTs with increases in GV and AC before the CT that are typical of BCTs. Therefore it is necessary to consider both the TSPs and the m(TSPs, t_m) when classifying CTs in voltage recordings of real seizure activity even if this restricts the amount of CTs that we can classify.

S11: SVM classifications and mean feature fit error vs. T

In this subsection we show how the proportion of filtered CTs classified as a particular CT type for a given T, denoted by $N_{\text{type}}(T)/N_{\text{filt}}(T^-, T^+) \in [0, 1]$, varies for $T \in [-10, 10]$ when setting $T^- = -20$ and $T^+ = 10$. These classification are provided by the SVM type-3 with $t_m = 8$ setup. The TSPs and m(TSPs, t_m) with $t_m = 8$ we use for this SVM are computed from the voltage recordings of rat K. CTs between the NS and S states are detected by our algorithm (using parameters in Table 3) and the CTs that are classified all satisfy C1-C5. In total there are $N_{\text{filt}}(-20, 10) = 95$ CTs eligible for classification, $\approx 10\%$ of the detected CTs. Despite only classifying such a small amount of the detected CTs, the proportion of CTs classified as a particular CT-type when $T = 2$ in this subsection are consistent with those presented in Table 4 for rat K ($\approx 25\%$ BCT, $\approx 25\%$ BNCT, $\approx 50\%$ NCT). We also assess the accuracy of these classifications in terms of how the TSPs and m(TSPs, t_m) of the classified voltage recordings compare to those of the model's output for each CT type.

SVM classifications: Figure S-17 (a) shows how $N_{\text{type}}(T)/N_{\text{filt}}(-20, 10)$ varies for different values of $T \in [-10, 10]$, this

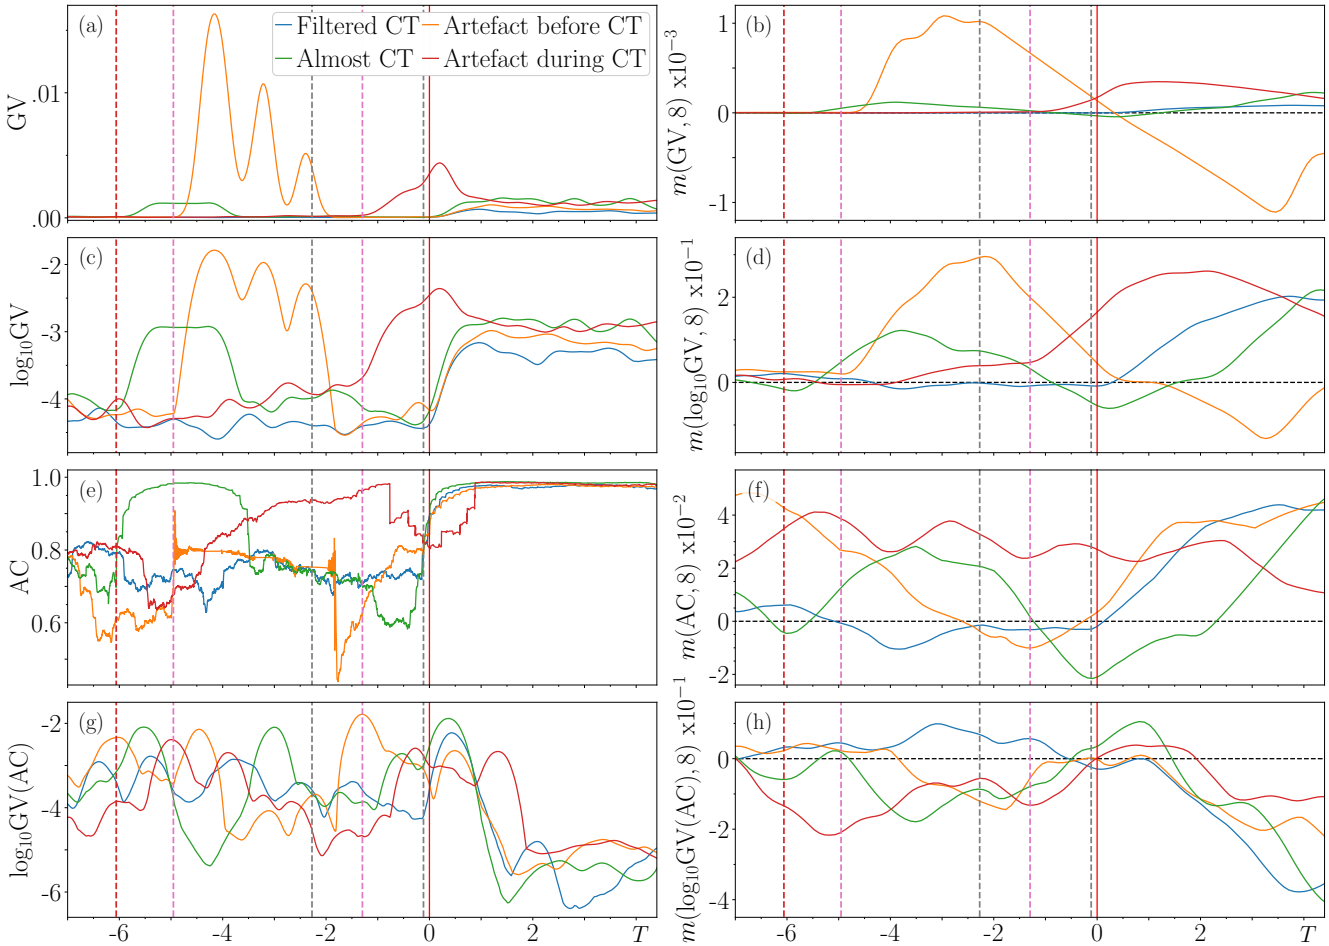

Figure S-16. TSPs and m(TSPs, t_m) of the voltage recording examples shown in Fig. S-15 and plotted versus the same T values where the colour of each curve corresponds to the different examples specified in plot legend. The vertical lines indicate (solid red) when $T = 0$, (dashed red) an almost-occurring CT from the NS to S state, (dashed pink and grey) the beginning and end of artefact activity, according to our algorithm.

tells us what proportion of the filtered CTs are classified as a particular CT type; BCT in blue, BNCT in green, and NCT in red.

Figure S-17 (a) shows that long before the CTs, i.e., for $T < -5$, $N_{\text{type}}(T)/N_{\text{filt}}(-20,10)$ is relatively similar for each CT type, i.e., the SVM cannot distinguish between CT types for $T < -5$. For $T \in [-5,0)$, Fig. S-17 (a) shows that $N_{\text{type}}(T)/N_{\text{filt}}(-20,10)$ varies significantly and there are times where different CT types appear to be the dominant seizure generation mechanism. However, just after the CT takes place, specifically, for $T \in [0,1]$, there is a major increase in the proportion of CTs classified as NCT, balanced by a significant decrease in the proportion of CTs classified as either BCT or BNCT. For $T \in (1,5]$ we note that $N_{\text{type}}(T)/N_{\text{filt}}(-20,10)$ remains relatively constant, thus providing a range of T values where more reliable classifications can be obtained. Similar to the results presented in Fig. 9, we see that for $T = 2$ in Fig. S-17 (a), NCTs are also the dominant seizure generation mechanism for this portion of rat K's CTs. For $T > 5$, the proportion of CTs classified as NCT remains relatively close to 0.5, BCT starts to increase and is balanced by a near equal decrease in BNCT. Fig. 6 shows this happens because the TSPs and m(TSPs, t_m) of each CT-type are much less distinguishable from each other for $T > 5$. Thus, the classification are less reliable for $T > 5$.

Mean feature fit error: While our main results are computed using an SVM that achieves perfect accuracy when classifying the model's CTs, we are not able to quantify the accuracy of the SVM when used to classify CTs in the voltage recordings in the same way since their generation mechanisms are unknown. Instead, we take the following steps to compare the classifications of CTs in the voltage recordings to the corresponding CTs in the model's output:

1. After classifying CTs in the voltage recordings, organise their features (i.e., the TSPs and m(TSPs, t_m)) into groups based on the classifications, i.e., groups of TSPs and m(TSPs, t_m) that correspond to BCTs, BNCTs, and NCTs.

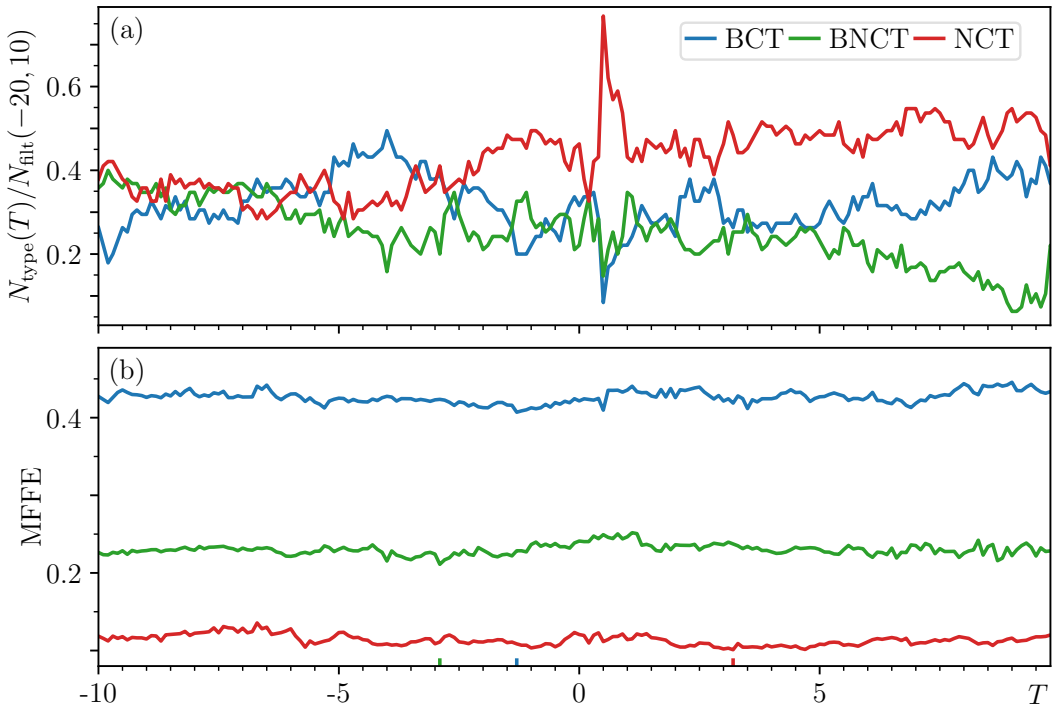

Figure S-17. Classification of seizure generation mechanisms and MFFE (mean feature fit error) versus T. (a) shows $N_{\text{type}}(T)/N_{\text{filt}}(-20,10)$ versus T for rat K, i.e., proportion of filtered CTs that are classified as (blue) BCT, (green) BNCT, or (red) NCT, for a given T with $T^- = -20$ and $T^+ = 10$. (b) shows the MFFE computed for each CT-type versus T, inward coloured ticks indicate the smallest MFFE for the corresponding CT-type.

2. Compute the mean value of each feature in the different groups for a given T, e.g., compute the mean GV at a given T for the group of CTs classified as BCTs, BNCTs, and NCTs and repeat for each feature.

3. Compare each mean to the equivalent quantities computed from the model's output (first and third columns of Fig. 6) by:

   (i) Performing a min-max normalisation of the means of each feature.

   (ii) Computing the root mean square error (RMSE) between corresponding normalised curves obtained from (i), see Figs. S-18 and S-19 for an example.

   (iii) Computing the 'mean feature fit error' (MFFE): the mean of the RMSEs over all features for a given T.

In Fig. S-17 (b) we plot the MFFE versus T having applied steps 1-3 listed above to the data used to generate Fig. S-17 (a). Figure S-17 (b) shows there is a relatively strong agreement between NCTs in the model and CTs in the voltage recordings that are classified as NCT. Figure S-17 (b) also shows there is an apparent weaker agreement between CTs that are classified as BNCT and an even weaker agreement for BCT.

To more closely investigate how well the features from the classified CTs in the voltage recordings mimic their counterparts in the model's output, we choose $T = -1.3$ as a point to visually compare the normalised means of the features from the model and the voltage recordings in Figs. S-18 and S-19. This time corresponds to when the MFFE is nearest to 0 for the BCT case.

Figures S-18 and S-19 show there are several qualitative similarities between the classified CTs in the voltage recordings and the corresponding CTs generated by the model. For instance, the first column in both Figs. S-18 and S-19 show that CTs classified as BCTs exhibit strong signatures of critical slowing down (CSD), i.e., increases in GV and AC before the CT, thus indicating that the rat's brain has drifted towards and then past a bifurcation point that generates the seizure.

Figure S-18. Comparing normalised values of TSPs of (dashed curves) the model to the corresponding quantities obtained from the (solid curves) voltage recordings that were classified by the SVM for $T = -1.3$ (the value of T where the best fit of the BCTs occur, shown in Fig. S-17 (b)). Each curve is coloured according to the class specified in the top left corner of panels (a)-(c). RMSE between the curves is quoted beside each pair of curves.

Figures S-18 and S-19 also show that, in comparison to the BCT case, there are significantly smaller RMSEs between the curves in the BNCT and NCT cases. For instance, panels (e) and (f) in Figs. S-18 and S-19 illustrate the strong agreement between the model and the voltage recordings in terms of the normalised $\log_{10}$GV and $m(\log_{10}\text{GV}, 8)$, further indicated by the small value of the RMSE quoted beside these curves. However, the same cannot be said for all features. For instance, panel (h) in Figs. S-18 and S-19 shows there is much poorer agreement between BNCTs in terms of the normalised AC and $m(\text{AC}, 8)$ curves. The large RMSE quoted beside these curves provides the greatest contribution to the MFFE shown in Fig. S-17 (b). The same issue arises for the BCT case as panel (g) in Figs. S-18 and S-19 show the RMSE for the normalised AC and $m(\text{AC}, 8)$ curves is also significantly large. What is important in both BCT and BNCT cases is that there is a clear increase in AC before the CT, however, it is the rate of increase that differs from the model's. Importantly, panel (i) in Figs. S-18 and S-19 show no gradual increase in AC before NCTs in the voltage recordings, which is consistent with the corresponding CTs generated by the model.

Figure S-20 shows the TSPs for a single example of each CT-type in the voltage recordings of rat K which were classified with the same SVM as in Figs. S-18 and S-19. In these examples, one can see a gradual increase in variance and autocorrelation prior to the BCT and BNCT, and no such increase in variance and autocorrelation prior to the NCT. When compared to Fig. S-11, one can observe in greater detail how TSPs corresponding to different CT-types in the voltage recordings align with those of the

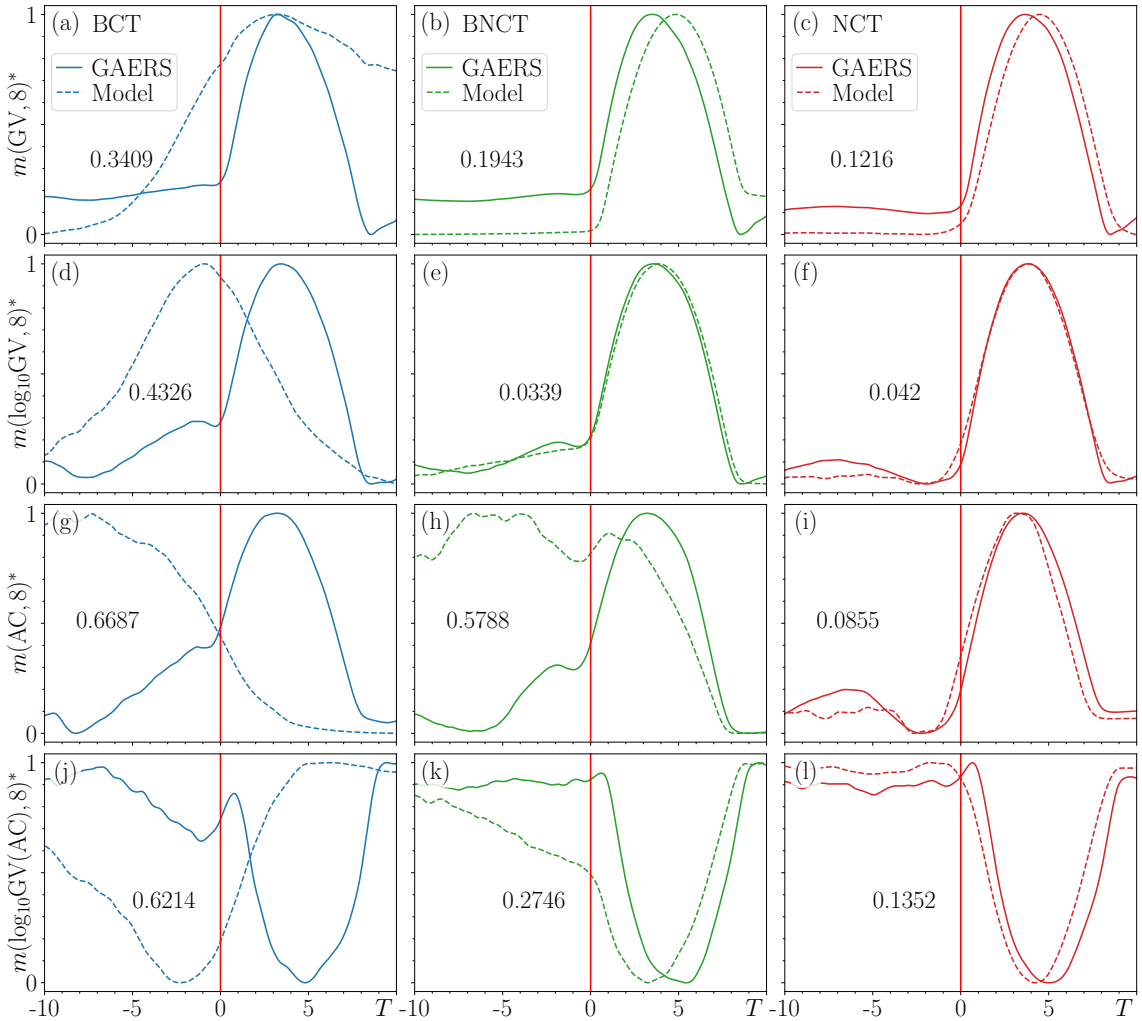

Figure S-19. Comparing normalised values of m(TSPs, t_m) with $t_m = 8$ of the (dashed curves) model to the corresponding quantities obtained from (solid curves) the voltage recordings that were classified by the SVM for $T = -1.3$ (the value of T where the best fit of the BCTs occur, shown in Fig. S-17 (b)). Each curve is coloured according to the class specified in the top left corner of panels (a)-(c). RMSE between the curves is quoted beside each pair of curves.

model's output.

S12: Applying filtering conditions to detected CTs

In this subsection we provide a more detailed account on how the conditions C1-C5 contribute to the information presented in Figs. 9 (d)-(f). More specifically, we compare the proportion of detected CTs that satisfy C1-C3 to those that satisfy C1-C5.

In Fig. S-21 we plot the proportion of detected CTs versus T^- for each rat for detected CTs that satisfy (i) C1-C3 (in blue) and (ii) C1-C5 (in black). From Fig. S-21 it is clear that rat K is the most affected by C4 and C5, followed closely by rat T. Rat S is least affected.

From further calculations we found that by changing $T^- =$ from -8 to $-\tau_{\mathrm{NS}} = -3$, the proportion of CTs that satisfy C1-C3 increased from $75-85\%$ to $99.7-100\%$. By imposing C4 and C5 we found that these proportions of detected CTs fell to 80% for rat S and 66% for rats T and K.

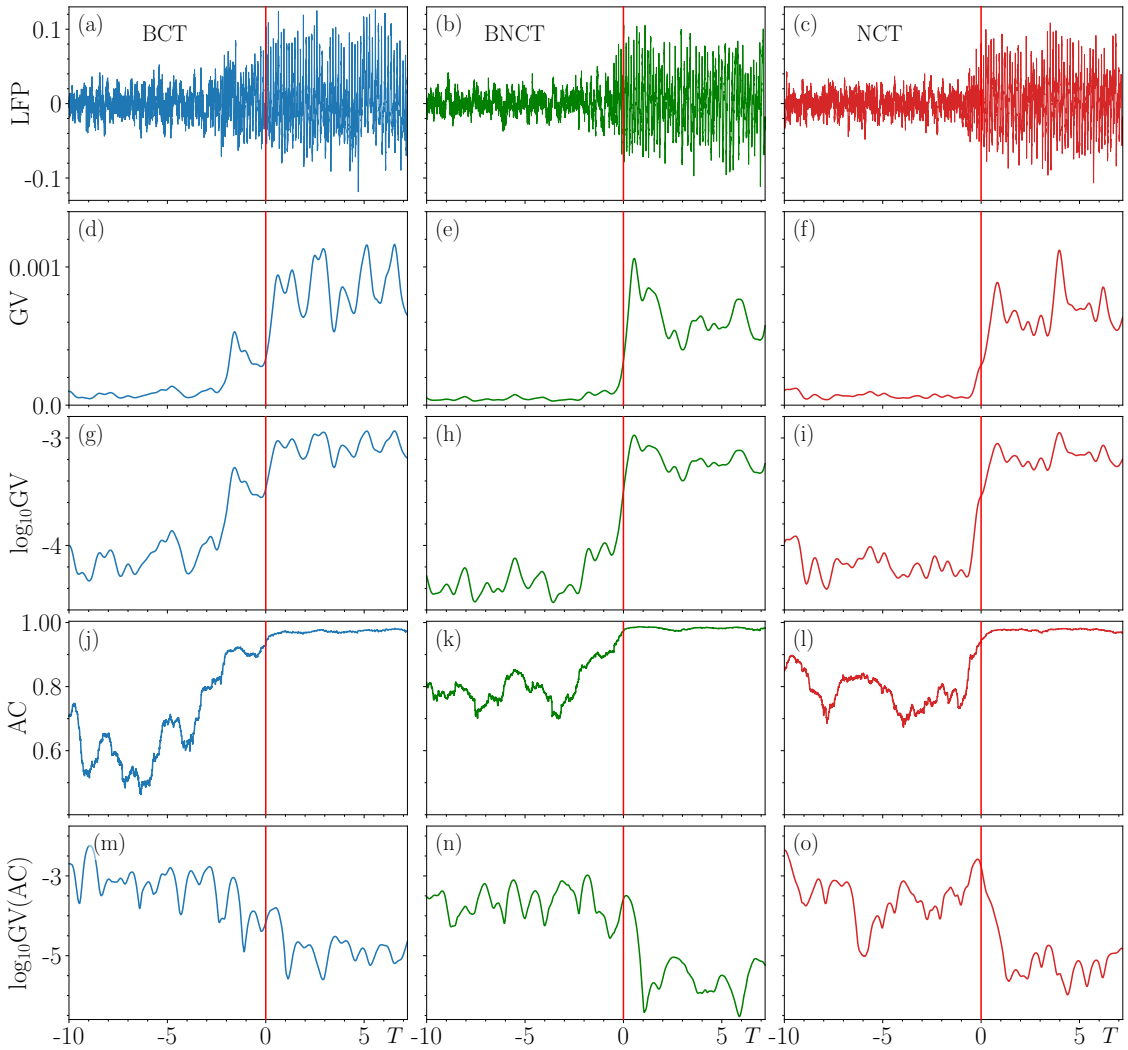

Figure S-20. Time series properties (TSPs) of voltage recordings for rat K near an example of a (left-hand column) BCT, (middle column) BNCT, and (right-hand column) NCT as classified by the same model-trained SVM used to generate Figs. S-18 and S-19. For each CT-type, the top row shows the LFP versus $T = t - t_1$ with t_1 detected by the algorithm. The remaining rows show how each of the TSPs, specified in the vertical axes labels and in Table 5, behave for the same values of T. Vertical red line in each panel emphasises when $T = 0$.

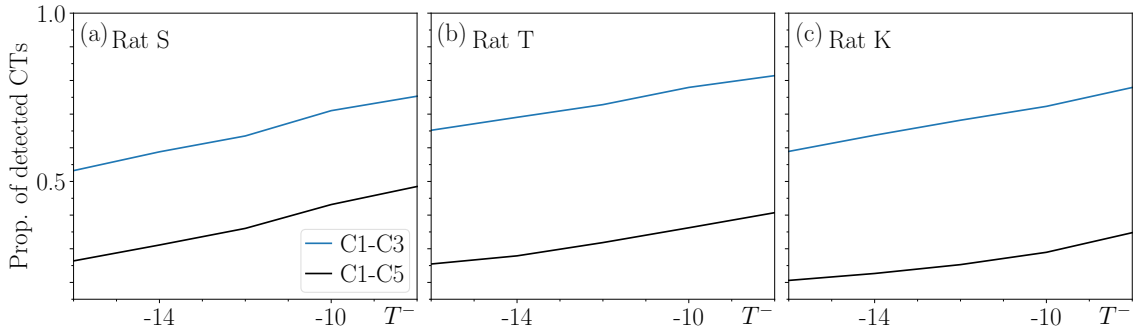

Figure S-21. Complimentary figure to Figs. 9 (d)-(f) showing the influence of artefacts and almost-occurring CTs happening nearby the CT when filtering the detected CTs of (a) Rat S, (b) Rat T, and (c) Rat K. Each panel shows the proportion of detected CTs that satisfy conditions (blue) C1-C3 and (black) C1-C5, i.e., proportion of detected CTs with/without artefacts and almost-occurring CTs happening nearby the CT.