

A framework for chaotic signature classification: Insights from Chua's circuit

Original

A framework for chaotic signature classification: Insights from Chua's circuit / Becchi, S., Spinazzola, E., Haliuk, S., Corinto, F., Chua, L.O., Pareschi, F., Vovchuk, D., Secco, J.. - In: CHAOS, SOLITONS AND FRACTALS. - ISSN 0960-0779. - 210:(2026). [10.1016/j.chaos.2026.118615]

Availability:

This version is available at: 11583/3012128 since: 2026-06-16T16:18:03Z

Publisher:

Elsevier

Published


DOI:10.1016/j.chaos.2026.118615

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



A framework for chaotic signature classification: Insights from Chua's circuit

Sara Becchi ^a,*^{*}, Elisabetta Spinazzola ^a, Serhii Haliuk ^{b,c}, Fernando Corinto ^a, Leon O. Chua ^d, Fabio Pareschi ^a, Dmytro Vovchuk ^b, Jacopo Secco ^a

^a Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi 24, Turin, 10129, Italy

^b Institute of Telecommunications, Riga Technical University, Kipsalas iela 6a, Kurzemes rajons, Riga, 1048, Latvia

^c Department of Radioengineering and Information Security, Yuriy Fedkovych Chernivtsi National University, Kotsyubyns'koho St, 2, Chernivtsi, 58002, Ukraine

^d Department of Electrical Engineering and Computer Sciences (EECS), University of California, Berkeley, 94720, CA, USA

ARTICLE INFO

Keywords:

Chaotic signature
Chua's circuit
Short-Time Fourier Transform (STFT)
Deep learning classification
Grad-CAM
Time-frequency representation

ABSTRACT

In this preliminary work we investigate the possibility to extract, from the analysis of a chaotic signal, a signature of the circuit that generated it. It is widely known that even small deviations in parameters of a chaotic circuit are reflected into generated signals with different dynamics and characteristics, not detectable by simple inspection and impossible to reproduce in another circuit. Being able to identify these differences, and being also able to verify if a sequence has been effectively generated by a given circuit will pave the way to new applications of chaotic circuits including, but not limiting to, security applications. In this paper we consider Chua's circuit, that is one of the simplest known chaotic circuits, and we show that by using deep learning and time-frequency transformations it is possible to extract a signature from chaotic sequences. Results are provided using two simulated datasets, and also two experimental datasets acquired from six hardware realizations of Chua's circuit. All signals are transformed into Short-Time Fourier Transform and classified with Convolutional Neural Network (CNN) models, achieving 98% macro-averaged balanced accuracy across the evaluated test sets, including 224 distinct chaotic parameterizations, unveiling a data-driven route to fingerprint chaos. As additional results, gradient-based interpretability and Permutational Multivariate Analysis of Variance (PERMANOVA) based statistical analysis are provided to support distinctive time-frequency features of the signatures.

1. Introduction

Chaotic circuits have been used in recent years in many field of Information Technology, in particular in cryptography. The unpredictability properties have been found to be appealing in many applications such as secure communication (including text and image encryption [1]) and random number generation [2–4]. Chua's circuit, introduced in 1983, remains a canonical example due to its ability to generate a remarkable diversity of chaotic attractors, despite its relatively simple hardware implementation [5]. The most interesting feature of chaotic signals is their unpredictability. Even if chaotic evolution arises from well-defined equations and is reproducible under ideal conditions, they are affected by a strong sensitivity to parameter variations and initial conditions. It is widely accepted that when considering noise and implementation errors, unavoidable in real-world circuits, we have a twofold

* Corresponding author.

E-mail address: sara.becchi@polito.it (S. Becchi).

<https://doi.org/10.1016/j.chaos.2026.118615>

Received 28 January 2026; Received in revised form 18 May 2026; Accepted 1 June 2026

Available online 9 June 2026

0960-0779/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).




Fig. 1. Chua's Circuit architecture. On the left is shown the circuit schematic and measurement nodes V_1 and V_2 . On the right, the piecewise-linear characteristic $f(x)$ of Chua's diode used in this work is depicted.

consequence: (i) it is impossible for the same chaotic circuit to generate the same evolution twice; (ii) it is impossible to find two chaotic circuits equal to each other.

While the first property is the most commonly used, here we want to exploit the second one in order to see, given a chaotic signal, if it is possible to identify the circuit that generated it. In other words, the aim of this paper is to see if a sort of "signature" exist in chaotic circuit, and more important, if this signature can be identified by only looking at the signals generated by the circuit. In particular, we aim to look for a set of stable, discriminative time-frequency features that recur in all generated signals, and that encode the system dynamic identity, remaining reproducible within a specific regime and measurably linked to both the topology and the specific parameter set of the chaotic circuit.

In security related applications this concept is widely used, for example in the physical identification of devices. Physical Unclonable Functions (PUFs) exploit intrinsic fabrication variations to generate unique, hard-to-clone responses, employed for secure device authentication [6]. Similarly, Radio Frequency (RF) fingerprint techniques identify a transmitter from hardware specific impairments [7].

To explore our idea, we focus on different implementation of the same Chua's circuits, analyzing how variations in its control parameters affect the resulting time-frequency patterns. The choice of the Chua's circuit is mainly based on its simplicity and its diffusion. Furthermore, thanks to a small set of tunable parameters, it offers a controllable platform bridging mathematical models and experimental reproducibility [8–10].

Since traditional signal-processing descriptors often fail in the identification of stochastic variability, we opted for a deep learning based solution. Deep learning is widely used as a practical framework for chaotic signal analysis as it may learn hierarchical, high level features, thus enabling robust discrimination of input data [11,12]. Furthermore, we obtain good results by introducing time-frequency transformation of chaotic signal. Time-frequency transformation outputs are encoded in a 2D color plot, that is then fed into a Convolutional Neural Network (CNN), especially designed for recognizing spatial patterns [13,14]. Recent studies by Uzun et al. and Akmeşe *et al.* show that image-based encodings enable high classification accuracies among canonical systems like Lorenz, Chen, and Rössler [15,16]. This strategy eliminates the need for manual feature extraction or the construction of complex theoretical models, allowing CNNs to learn the distinctive features directly from the visual representations of the signals.

The You Only Look Once (YOLO) architecture has been chosen for the implementation of the CNN. This architecture is widely employed for analyzing spectrogram images of complex RF signals [17], audio signal analysis through mel-spectrograms [18] and real-time telecommunication network classification [19]. The choice of using YOLO, and specifically YOLOv11, is supported by several applied studies confirming the network's ability to successfully analyze data transformed in complex image representations. In the field of signal processing, SPEC-YOLO [19] has demonstrated outstanding performance in recognizing wireless signals converted into spectrograms, while in audio processing, the same approach has been used to detect deepfake anomalies [18].

In our analysis the spectrograms are plotted applying the Short-Time Fourier Transform (STFT) to signals generated by the Chua's circuit. The computation of the squared magnitude of the STFT yields a 2-D map that can be processed directly with image-based deep learning methods.

To further interpret the classification results, the Grad-CAM AI explainability technique is employed, in order to highlight the regions in the spectrogram where the network focuses its attention when identifying which circuit configuration generated a given time-frequency pattern. The salient points extracted through this method are subsequently subjected to PERMANOVA statistical analysis to confirm that the features identified differ significantly between different chaotic regimes.

The paper is organized as follows. Section 2 provides an introduction on Chua's circuit. Section 3 describes a first, simple case study, while Section 4 presents results of a more complex and more realistic case study with more chaotic classes. Extended insights, including the use of Grad-CAM to explore class-dependent time-frequency patterns in the spectrograms have been reported in Section 5, with a statistical validation of the identified salient regions. In the same section, additional considerations on the combination of chaotic classes are done. Finally, the analysis of real hardware Chua's circuits as a final validation step is proposed in Section 6. The overall results are then discussed in Section 7, followed by a brief conclusion.

2. Chua's circuit theory

The Chua's circuit was selected as a reference model, as it represents one of the most extensively studied and versatile chaotic electronic circuit. Its schematic is shown in Fig. 1.

We adopt the normalized mathematical formulation of Chua's system [9,20], described by three first-order differential equations (1) governing the evolution of the state variables:

$$\begin{cases} \frac{dx}{dt} = \alpha \cdot (y - x - f(x)) \\ \frac{dy}{dt} = x - y + z \\ \frac{dz}{dt} = -\beta \cdot y - \gamma \cdot z \end{cases} \quad (1)$$

where x and y correspond to the capacitor voltages (referred to as V_1 and V_2 in Fig. 1), and z represents the current through the inductor L . The parameters α , β , and γ arise from the linear RLC sub-circuit and the adopted normalization of time and state variables. Unlike simplified lossless formulations (often obtained by neglecting parasitic losses), we explicitly retain the γ term in the third state equation of (1) to account for linear damping (resistive effects in the inductive branch), which is relevant when aligning simulated trajectories with hardware realizations [21]. The nonlinear term $f(x)$ implements the so-called Chua's non-linear diode. It is shown in Fig. 1, and described as:

$$f(x) = m_2x + \frac{1}{2} \left[(m_0 - m_1)(|x + a_1| - |x - a_1|) + (m_1 - m_2)(|x + a_2| - |x - a_2|) \right]. \quad (2)$$

The non-linear function $f(x)$ consists of three segments with slopes (m_0, m_1, m_2) and breakpoints (a_1, a_2) , and it is the source of the circuit's dynamical richness. The relationship between the hardware parameters of the circuit and those used in the mathematical model can be found in [22].

We define the "signature" of a Chua's circuit as the set of normalized parameters $(\alpha, \beta, \gamma, m_0, m_1, m_2, a_1, a_2)$. We will consider from now on that two Chua's circuits are different if each of them is identified by a different signature, i.e., a different set of parameters, being this obtained either by intentionally modifying parameters at a circuital level, or by considering the tolerance of passive devices (resistors, capacitors, inductors) in a real implementation.

Variations in the set of parameters are directly reflected in the time-frequency representation of the generated signals. Due to the chaotic nature of the Chua's circuit, even small variations in the set of parameters are reflected in large variations in the time-frequency properties of the generated signal. This property should allow distinguishing between two signals generated by two Chua's circuits with different signatures. On the contrary, the chaotic nature of the Chua's circuit prevents the reconstruction of the signature from the observation of the time-frequency properties of the generated signals. This is aligned with the fundamental properties of a signature to not be inferred, estimated, or reverse-engineered from simple observation.

As a final comment, we would like to clarify a few aspects that, given the preliminary nature of this study on the signature of chaotic circuits, we deliberately leave outside the scope of this manuscript so as not to increase its complexity. The first one is noise. Real implementations of any chaotic circuit will inevitably be affected by noise, which will alter the time-frequency properties of the generated signal. Adopting the simplest and most common model of additive white Gaussian noise would not impair the identification of the signature, as it would affect all spectra in the same way, without altering the relative differences of time-frequency structures that are essential for distinguishing one signature from another one. Instead, adopting a more complex, signal-dependent noise model would make the problem analytically intractable. The second aspect is temperature. A change in temperature is expected to introduce variations in the time-frequency properties of all generated signals. However, accounting for this effect would require a detailed model of the temperature sensitivity of all parameters, which lies beyond the scope of this preliminary study. Finally, an important property of any signature-based system is its resilience to linear transformations (e.g., filtering). A thorough investigation of this aspect is left for future work.

Nevertheless, the measurements presented in Section 6 were obtained in an open, unshielded environment where realistic noise sources are present, and without any control of temperature. The results obtained corroborate the hypothesis that noise and temperature variations are not critical issues for this study.

3. A simple case study

A first and very simple case study has been designed to investigate if a CNN classifier is capable to identify the signature for a small number of Chua's circuit.

In order to build a dataset for training and validating the CNN, a dataset containing four classes of signal is built. Three of them are coming from three different realizations of Chua's circuit; the fourth one is different, and created to represent the "none of the above" case, i.e., when no match with a known circuit is found. This class has been added to prevent a limitation of CNN classifiers, that have to provide a match with one of the class in any case.

This case study has been considered by means of numerical simulations. The normalized equations in (1) are integrated using the Runge-Kutta method of order 5 [23]. The simulated trajectories for variables x and y are extracted, in the aim of emulating a real circuit where only the two voltages V_1 and V_2 could be easily measured, whereas getting the current of the inductor I_L is more difficult.

Hereinafter, for simplicity, we will denote each class as S_s^p , where $p \in \{1, 2, 3, 4\}$ represents the circuit parameters set of each class, and s denotes the type of signal output considered between x and y .

In detail, linear parameters for all classes are fixed to $\alpha = 9.237$, $\beta = 16.8$ and $\gamma = 0.0042$ [9]. The corresponding nonlinearity parameters are different, and shown in Table 1.

Table 1

Nonlinearity parameters $(m_0, m_1, m_2, a_1, a_2)$ used in the first case study chaotic models. All simulations were performed with $\alpha = 9.237$, $\beta = 16.8$, and $\gamma = 0.0042$. Parameters for $S^4_{\{x|y\}}$ were independently sampled from continuous uniform distributions over the illustrated ranges below.

Class	m_0	m_1	m_2	a_1	a_2
$S^1_{\{x y\}}$	-1.238	-0.6665	500.0	1.00	6.88
$S^2_{\{x y\}}$	-1.230	-0.6900	440.0	1.10	7.15
$S^3_{\{x y\}}$	-1.200	-0.7200	300.0	1.12	7.00
$S^4_{\{x y\}}$	Random	Random	Random	Random	Random
Ranges ^a	[-1.8, -0.7]	[-1.0, -0.3]	[200, 700]	[0.8, 1.5]	[6.0, 8.5]

^a Allowed intervals for random parameters.

Within each of the classes $S^1_{\{x|y\}}$ to $S^3_{\{x|y\}}$, only the initial conditions are changed, so all signals share the same underlying chaotic regime and differ only by their trajectory on the same attractor.

The construction of the fourth class $S^4_{\{x|y\}}$ is different. To implement a “generic” class, the nonlinearity parameters are not fixed but randomly sampled within broad chaos-preserving ranges that overlap the first three regimes. As a result, $S^4_{\{x|y\}}$ consists of multiple chaotic regimes with nearby, yet distinct, parameter settings, forming a heterogeneous class that does not correspond to any single reference configuration. Trajectories in $S^4_{\{x|y\}}$ resemble those of $S^1_{\{x|y\}}$ - $S^3_{\{x|y\}}$ but never match any of their exact parameter sets.

In the latter case, to ensure that all the generated solutions remained within a chaotic regime, the full Lyapunov spectrum $(\lambda_1, \lambda_2, \lambda_3)$ of the normalized Chua system was computed for each case using a standard Benettin-type algorithm with tangent vector renormalization [24]. This method quantifies the average exponential divergence of initially nearby trajectories by integrating the associated Jacobian dynamics and periodically rescaling perturbation vectors. The finite-time estimator of each Lyapunov exponent over K QR steps of durations Δt_k is computed as follows [25]:

$$\hat{\lambda}_i(T) = \frac{1}{T} \sum_{k=1}^K \ln r_{ii}^{(k)}, \quad T = \sum_{k=1}^K \Delta t_k, \quad i = 1, \dots, n. \tag{3}$$

As expected for a 3-D dissipative flow, chaotic regimes exhibit $\lambda_1 > 0$, $\lambda_2 \approx 0$, $\lambda_3 < 0$, and $\sum_i \lambda_i < 0$ [26].

In detail, after discarding an initial transient, the Largest Lyapunov Exponent (LLE) was estimated over long integration times with strict numerical tolerances. For downstream selection we retain only the sets of parameters that exhibit $LLE > 0.05 \text{ s}^{-1}$. While a positive maximal Lyapunov exponent is the standard operational marker of deterministic chaos [27,28], finite-time estimation and numerical effects can yield spuriously small positive values near periodic or quasi-periodic windows, this criterion is advisable to avoid borderline cases [29,30].

The initial state vector $\{x_0, y_0, z_0\} \in \mathbb{R}^3$ was initialized such that each component was independently drawn from a uniform distribution, i.e., $\{x_0, y_0, z_0\} \sim \mathcal{U}([-1, 1]^3)$. After a burn-in period of 50 s, signals were sampled for 30 s at a frequency of 300 Hz. Preliminary Welch power-spectral analysis showed that at least 95% of the signal power is concentrated below 125 Hz, with negligible energy at higher frequencies. Choosing 300 Hz as sample frequency therefore satisfies the Nyquist criterion, ensuring that the relevant chaotic dynamics are captured without aliasing, while avoiding unnecessarily dense sampling that would inflate the dataset size and the cost of time-frequency computations. For each class we generated 3000 independent samples changing initial condition vectors.

An example of the phase portraits for the variables x - y for each of the four classes is depicted in Fig. 2.

To convert the signals into interpretable patterns, we employed the Short-Time Fourier Transform (STFT) and derived log power spectrogram images from it. This method provides the exploration of time-varying spectral representation and allows the examination of signal dynamics along both the time and frequency dimensions. Thanks to its windowed formulation, the STFT can be applied to nonstationary signals such as chaotic trajectories [31]. We adopt the STFT because it offers a strong trade-off between representation quality and computational cost and, unlike bilinear distributions, as Wigner-Ville, does not suffer from cross-term interference that obscures multicomponent chaotic signals [32]. The equation for the computation of the transform is given by

$$X[m, k] = \sum_{n=0}^{N-1} x[n + mR] w[n] e^{-j2\pi kn/N} \tag{4}$$

where $x[n]$ is the discrete-time signal sampled at f_s and $w[n]$ is the $N = 1024$ -sample Hamming window applied to each segment to reduce spectral leakage [33]. A 50% overlap between each frame is adopted [34]. The spectrogram images are obtained by mapping the STFT magnitude, expressed in dB, to a color scale. The logarithmic representation is fundamental to compress dynamic range, turn multiplicative gain differences into additive offsets, facilitating comparison across samples and enhancing low-energy structures visibility with a fixed color map. The phase component of the complex STFT was not exploited, as log magnitude-based spectrograms proved to be sufficient for our classification goal.

Fig. 2 includes an example of the log power spectrogram images, for the four classes (in particular, for the spectrogram associated to the x variables given from the corresponding phase portraits).

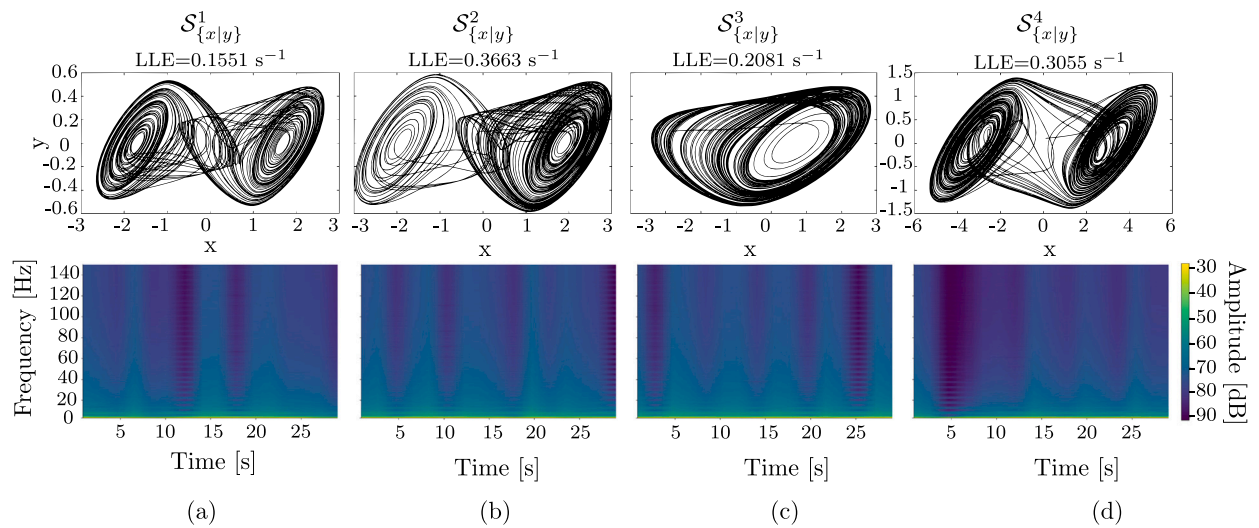


Fig. 2. Example of phase portraits x - y (top figures) and log power spectrogram images derived from the x output (bottom) of for the four considered classes for the first case study. Panels (a-d) passed the quality checks (bounded trajectories, LLE 0.05 s^{-1}) and exhibit the expected double-/single-scroll morphology after transient removal.

For the automatic classification of spectrograms, a YOLOv11-Medium was used as classification network. YOLO11m-cl's architecture is configured with a classification head that outputs class probabilities using softmax activation [35] and has been optimized to reduce computational cost without compromising accuracy performance and model versatility.

Two distinct YOLO-models were trained, one classifier was trained on spectrograms derived from x signals and one on spectrograms derived from y .

Each model was trained using a transfer learning approach, starting from weights pre-trained on the ImageNet Dataset, ensuring generalization and preventing overfitting [36]. The models were trained for 200 epochs with an early stopping criterion (patience = 5) and automatic mixed-precision training (AMP) to improve efficiency. In practice, early stopping was never triggered and the full 200 epochs were completed in all training runs. The value of 200 epochs was chosen as a compromise between allowing the networks to fully converge and limiting the overall computational cost across the four training procedures.

All datasets were divided into training (60% of input dataset), validation (20%), and test set (20%) subsets using a random stratified sampling strategy to ensure that the representativeness of each class was maintained across the dataset split. This split ratio provides a balanced trade-off between ensuring sufficient data for model learning and retaining independent samples for robust validation and final performance evaluation. Model performance was monitored on the validation set metrics of loss and accuracy during training process, while the final evaluation was carried out on the held-out test sets using accuracy and confusion matrices. Classification results are summarized in Fig. 3, depicting the confusion matrix achieved for the classification of the two considered circuit outputs x and y . Most samples lie on the main diagonal, indicating only few misclassifications.

Performances were assessed choosing metrics of Balanced Accuracy, Recall and Precision, defined class-by-class according to the value of True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN) as

$$\text{BalancedAccuracy} = \frac{1}{2} \left(\frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{TN} + \text{FP}} \right); \quad (5)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}; \quad (6)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}. \quad (7)$$

Table 2 shows performance metrics per class for this Case Study. Macro-level scores (obtained by averaging the per-class metrics across all classes) show that, for the x output, the model achieved a macro balanced accuracy of 99.14%, with a macro precision of 98.72% and a macro recall of 98.71%. For the y output, the macro balanced accuracy was 99.47%, with macro precision and macro recall both equal to 99.12%.

4. A more complex case study

Given the promising results of the simple case study, a more complex and more realistic case study has been considered. This time, a dataset is constructed by simulating 200 different Chua's circuits obtained by randomly varying both the nonlinearity coefficients (m_0, m_1, m_2, a_1, a_2) and the global parameters α , β , and γ as in Table 3. As in the previous case study, for each configuration the

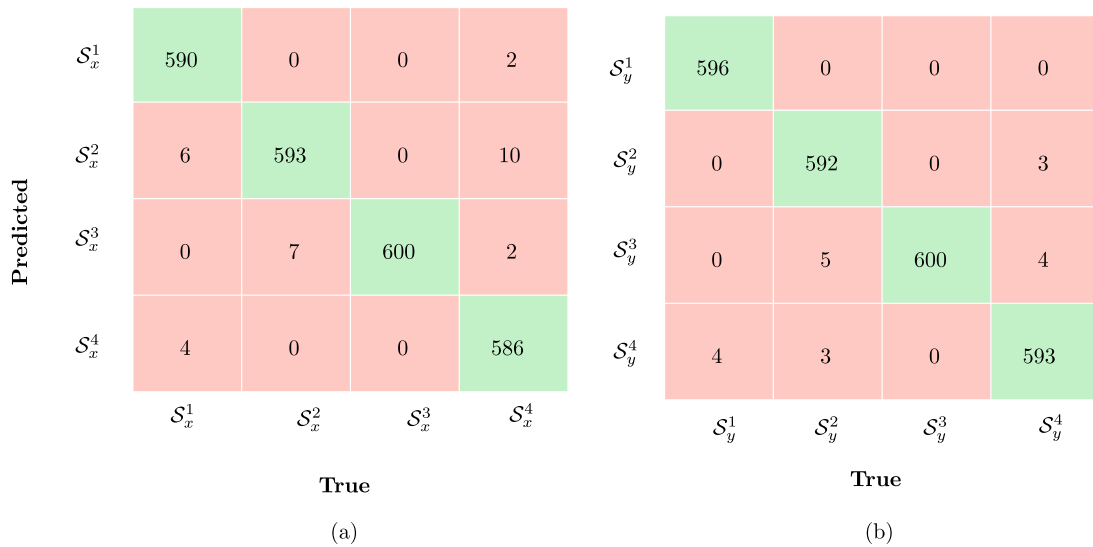


Fig. 3. Confusion matrix (a) obtained on the test set for the x signals and (b) for the y signals of first Case Study.

Table 2
Classification performance metrics for signals in the first case study.

Class	Samples	Precision	Recall	Balanced accuracy
S_x^1	600	99.66%	98.33%	99.11%
S_x^2	600	97.37%	98.83%	98.97%
S_x^3	600	98.52%	100%	99.75%
S_x^4	600	99.32%	97.67%	98.72%
S_y^1	600	100%	99.33%	99.67%
S_y^2	600	99.50%	98.67%	99.25%
S_y^3	600	98.52%	100%	99.75%
S_y^4	600	98.83%	98.83%	99.22%

Table 3
Sampling ranges for parameters of second Case Study. All parameters were sampled uniformly within the ranges specified below.

Parameter	Range
α	[8.000, 15.000]
β	[10.000, 25.000]
γ	[0.00000, 0.010000]
m_0	[-1.8, -0.7]
m_1	[-1.0, -0.3]
m_2	[200, 700]
a_1	[0.8, 1.5]
a_2	[6.0, 8.5]

Lyapunov spectrum and the LLE was estimated to ensure that the selected combination of parameters led to chaotic behavior. Once a valid parameter set was identified, it defined one class. For each class 4000 signals were simulated for $T = 30$ s with different initial conditions for each simulated realization.

Given the very high number of classes, this dataset seeks to approximate a broad and representative subset of the chaotic dynamics that Chua’s circuit exhibits. By widening the admissible parameter ranges and randomly sampling within chaos-preserving intervals, we expose the classifier to many distinct regimes and verify whether the extracted fingerprints remain discriminative as the number of configurations increases. In this case, an additional class corresponding to a failed match with any of the other classes has not been considered.

Classification results when applying the same procedure discussed in Section 3 are as good as that of the simple case. Qualitative confusion matrices are shown in Fig. 4(a) for x and Fig. 4(b) for y , showing a pronounced diagonal trend of correct classifications with only a limited number of errors.

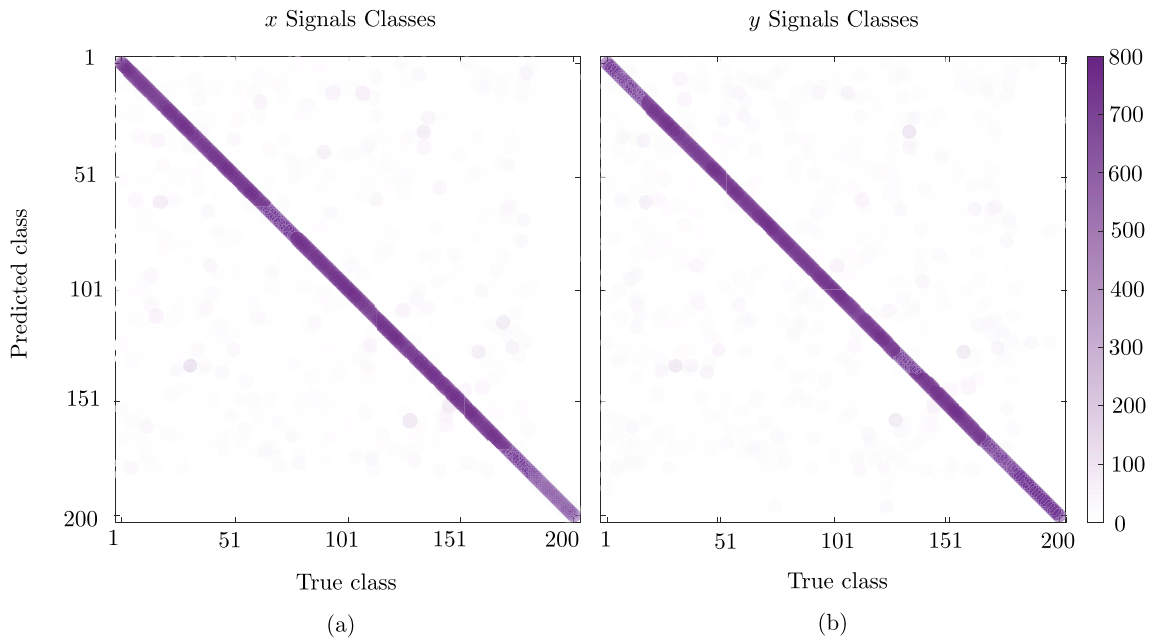


Fig. 4. Confusion matrices of the 200-class classifier on the test set ($N = 800$ signals per class) for the x (left) and y (right) signals. Axes report true (x-axis) and predicted (y-axis) classes. In both cases, the strong dominance of the main diagonal, with only a few faint off-diagonal entries, indicates that misclassifications are rare compared with the number of correctly classified samples.

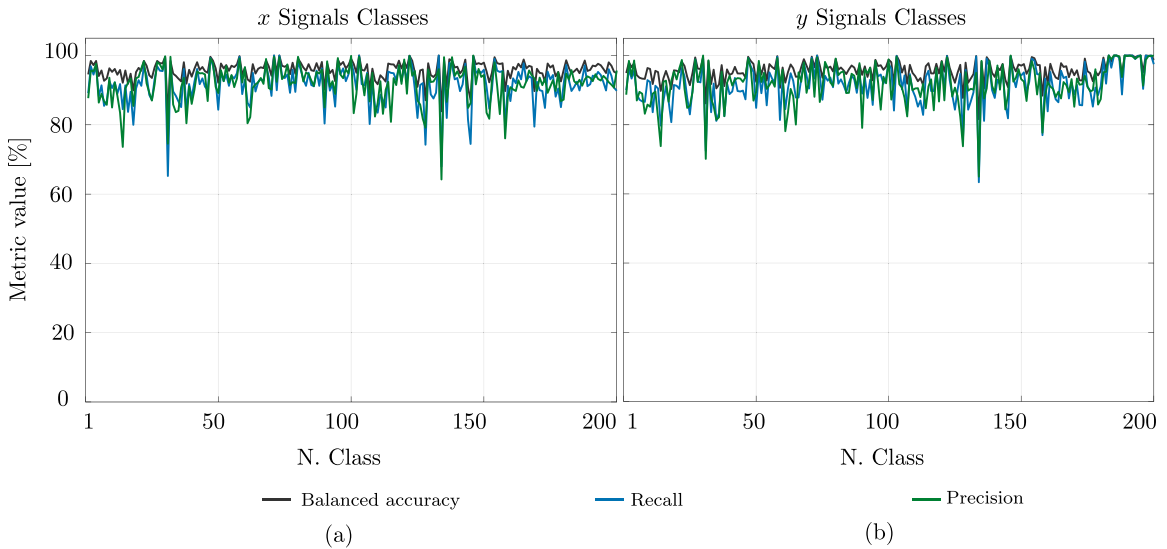


Fig. 5. Per-class performance over the 200 classes for x (left) and y (right) signals. Balanced accuracy, precision and recall computed on the test set are shown for each class.

To visualize the performance spread across the 200 classes, Fig. 5(a) and Fig. 5(b) show in detail the per-class balanced accuracy, precision and recall on the test sets for x and y , respectively. Overall, the curves remain consistently high across most classes, with only a limited number of dips that indicate a small subset of regimes that are harder to discriminate. Importantly, the two channels exhibit very similar trends and ranges, suggesting that the proposed pipeline extracts comparable discriminative information from both circuit outputs. In terms of macro results, this case study achieved a macro balanced accuracy of 95.82% on x , with macro precision and macro recall of 91.77% and 91.69%, respectively. Very similar performance was observed on y with macro balanced accuracy of 95.82%, macro precision of 91.80% and macro recall of 91.68%.




Fig. 6. Grad-CAM activation map for a representative test spectrogram. Warmer colors indicate time-frequency regions with higher contribution to the classification decision. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5. Extended insights

While the achieved classification results demonstrate strong discriminative performance, they do not by themselves explain what time-frequency structures the network exploits to separate chaotic regimes. Aim of this section is to provide some considerations based on the results of the first two case studies.

5.1. Explainability and statistical validation

Rather than employing explainable AI only as a post-hoc visualization tool, in this work we use Gradient-weighted Class Activation Mapping (Grad-CAM) [37] to extract class-specific time-frequency patterns that may act as spectral features of chaotic signatures. The key question is whether the most salient spectrogram regions identified by the network are (i) stable within each class and (ii) distinct across classes, thus revealing reproducible and discriminative structures.

Grad-CAM produces a class-discriminative saliency map $M_c(i, j)$ for class c by combining the feature maps of a convolutional layer with the gradients of the target score, without requiring architectural changes or retraining [37]. In image-based pipelines, including those operating on spectrograms, it has been shown to highlight stable time-frequency bands associated with specific classes, such as species-discriminative harmonics in bioacoustics [38], synthesis artifacts in audio signals [39], and modulation-dependent structures in wireless/RF signals [40]. These results motivate its use here not only to “explain” decisions, but to isolate candidate features of each chaotic regime in the time-frequency domain.

In our setting, Grad-CAM was applied to the last convolutional layer of the classification trained network, which captures high-level semantic features while retaining sufficient spatial resolution for localization. For statistical representativeness, Grad-CAM maps were computed for the Case Study II test sets only, since this scenario provides the broadest class coverage and the largest sample size (200 classes per channel), enabling a more robust and statistically meaningful multivariate analysis.

For each test spectrogram and each predicted class c , we obtained a saliency map $M_c(i, j)$, as shown in Fig. 6, and retained only the top 5% of its activation values. Specifically, a class-specific threshold τ_c was defined as the 95th percentile of all pixel intensities, computed across all maps of class c and a binary mask $B_c(i, j)$ was constructed as

$$B_c(i, j) = \begin{cases} 1 & \text{if } M_c(i, j) \geq \tau_c \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

thus retaining only the most salient time-frequency points contributing to the classification decision.

The corresponding indices (i, j) were then mapped to their associated frequency and time coordinates using linear scaling. This interpolation assigns to each pixel the center of its corresponding time-frequency bin.

For each chaotic class, the (f_i, t_j) coordinates selected by Grad-CAM were saved over all test samples, yielding an empirical distribution of salient points in time-frequency space. Subsequently, these distributions are analyzed statistically to test whether Grad-CAM has uncovered class-dependent features of the signatures rather than sample-specific artifacts.

In order to provide a statistical support to the “signature” hypothesis and complement the qualitative Grad-CAM inspection, for each chaotic class of second Case Study, we generated 800 signal realizations (considering x and y separately), computed the corresponding STFT spectrograms, and applied Grad-CAM. To statistically compare these class-wise saliency distributions, we employed Permutational Multivariate Analysis of Variance (PERMANOVA) [41]. The null hypothesis is H_0 : the Grad-CAM saliency patterns do not depend on the class, meaning that spectrograms from the same class are not more similar to each other than

Table 4

Classification outputs for the digitally summed Case Study I signals x and y (collectively denoted as $\{x | y\}$) in the fixed-regime classes $S^1_{\{x|y\}} - S^4_{\{x|y\}}$.

Input signal class	Output classification class				Accuracy
	S^1_x	S^2_x	S^3_x	S^4_x	
S^{1+2}_x	0	0	0	600	100%
S^{1+3}_x	0	0	0	600	100%
S^{2+3}_x	0	0	0	600	100%
S^{1+2+3}_x	0	0	0	600	100%
	S^1_y	S^2_y	S^3_y	S^4_y	
S^{1+2}_y	0	0	0	600	100%
S^{1+3}_y	0	0	0	600	100%
S^{2+3}_y	0	0	0	600	100%
S^{1+2+3}_y	0	0	0	600	100%

spectrograms from different classes (any apparent grouping is compatible with random label assignments). PERMANOVA tests this by comparing within-class and between-class dissimilarities and estimating the p -value via random permutations of the class labels.

In our case, the dissimilarity between two spectrograms was computed as the Hausdorff distance between their sets of salient points [42]. In simple terms, this distance measures the mismatch between two highlighted regions as follows: for each salient point in one distribution, we find the closest salient point in the other and the Hausdorff distance is the largest of these “closest-point” distances (and symmetrically vice versa). Therefore, it is small only when every highlighted region in one distribution lies close to some highlighted region in the other, and it becomes large if at least one salient region has no counterpart nearby.

We applied PERMANOVA to all classes separately for x and y . The analysis yielded $F = 35.35$ with $p = 0.001$ for x and $F = 27.50$ with $p = 0.001$ for y , leading to rejection of H_0 in both cases. Here, a large pseudo- F means that the average dissimilarity between classes is much larger than the average dissimilarity within the same class, and the very small p -value means that such a separation would be extremely unlikely if class labels were irrelevant. Practically, this supports our hypothesis: the Grad-CAM salient regions are consistent across realizations of the same circuit (stable within-class features) and systematically different across circuits (distinct between-class patterns), hence providing evidence of class-specific time-frequency features of the different signatures for both x and y .

5.2. Combination of multiple signals

One of the questions we need to answer to validate achieved classification results is what happens if we consider the combination of two signals coming from two known classes. With time-frequency properties intermediate between that of the two original classes, it is unclear if the combination signal will be recognized as a new signal, or as one of the two original signals.

To see this, we design the following experiment based on the first case study.

We considered three pairwise sums S^{1+2}_x , S^{1+3}_x and S^{2+3}_x , where with the notation S^{1+2}_x we intend the class of signals obtained as the sum of the x output from a circuit in class 1 and the x output from a circuit in class 2. A fourth combination is considered as the three-regime sum S^{1+2+3}_x , and other four classes are built in the same way for y signals.

Each new class consists of 600 samples and we have classified them according to the first case study fixed-regime classes $S^1_{\{x|y\}} - S^4_{\{x|y\}}$. The test has been conducted on the original CNN model *without re-training the network*. The results have been illustrated in Table 4, where the respective classifications for x and y signals are shown. The spectrograms of the combined signals are assigned to S^4_x , the class defined to capture inputs that do not correspond to any of the fixed classes.

In other words, combining two or more signals does not cause the classifier to confuse the mixture with any of the original single-regime classes ($S^1_{\{x|y\}} - S^3_{\{x|y\}}$). Instead, these mixtures give rise to a distinct new class, which is correctly assigned to $S^4_{\{x|y\}}$, explicitly designed as a heterogeneous “none of the above” category.

This is important, as suggests us that to increase the number of classes it is not necessary to change the parameters of a chaotic circuits, but it is also possible to combine in different ways existing chaotic circuits.

6. Final validation: an experimental case study

To experimentally validate the proposed method, we consider four independently operating hardware realizations of Chua’s circuit described in [9]. Each instance (denoted as $cc \in \{1, 2, 3, 4\}$) was assembled with different component values, leading to distinct (yet chaotic) dynamics due to tolerance-induced parameter spread. From each circuit, we acquired the voltages V_1 and V_2 at the two capacitor nodes (corresponding to the x and y state variables), with sampling frequency $f_s = 400$ kHz and duration $T = 5$ s per trace.

In order to increase the number of chaotic classes, given the observation in Section 5.2, we also considered linear combination (in detail, the sum achieved directly at hardware level) of multiple chaotic signals. Specifically, eight classes are considered for the V_1 channel (and eight classes for V_2 as well):




Fig. 7. Confusion matrix obtained on the test set (a) for the V_1 signals and (b) for the V_2 signals of the Experimental Dataset.

Table 5
Validation on real hardware Chua’s circuit.

Class	Samples	Precision	Recall	Balanced accuracy
$C_{V_1}^1$	10	100%	100%	100%
$C_{V_1}^2$	10	100%	100%	100%
$C_{V_1}^3$	10	83.3%	100%	98.57%
$C_{V_1}^4$	10	100%	100%	100%
$C_{V_1}^{1+2}$	10	100%	100%	100%
$C_{V_1}^{3+4}$	10	100%	80%	90%
$C_{V_1}^{1+2+3}$	10	100%	100%	100%
$C_{V_1}^{1+2+3+4}$	10	100%	100%	100%
$C_{V_2}^1$	10	100%	100%	100%
$C_{V_2}^2$	10	100%	100%	100%
$C_{V_2}^3$	10	100%	100%	100%
$C_{V_2}^4$	10	100%	100%	100%
$C_{V_2}^{1+2}$	10	100%	100%	100%
$C_{V_2}^{3+4}$	10	100%	100%	100%
$C_{V_2}^{1+2+3}$	10	100%	100%	100%
$C_{V_2}^{1+2+3+4}$	10	100%	100%	100%

1. **Single-circuit classes:** 4 classes of V_1 signals from $cc = 1$ to $cc = 4$;
2. **Two-circuits summation:** one class made up of the sum of V_1 signals of $cc = 1$ and $cc = 2$, and of $cc = 3$ and $cc = 4$;
3. **Three-circuits summation:** one class made up of the sum of V_1 signals of $cc = 1$, $cc = 2$ and $cc = 3$;
4. **Four-circuits summation:** one class made up of the sum of V_1 signals from $cc = 1$ to $cc = 4$.

Hereinafter, we will denote each of these classes as C_s^{cc} , where $cc \in \{1, 2, 3, 4\}$ indicates the circuit generating the considered signal, or all the circuits considered for generating the summation according to the notation of Section 5.2. The subscript s specifies the type of signal between V_1 and V_2 .

These signal are processed as in Section 3, with some adjustment due to the fact that signals have a different sampling rate and duration, and the available instances are limited.

In particular, before spectrogram computation the raw traces acquired at $f_s = 400$ kHz were downsampled to $f_s = 4$ kHz to reduce computational load while preserving the frequency content of interest. STFT was then applied to the downsampled signals (duration $T = 5$ s) using a window length $w_l = 0.909$ s with 50% overlap, yielding a frequency resolution $\Delta f = 1.10$ Hz and a temporal hop $\Delta t = 0.455$ s.

In the training, the early-stopping criterion terminated the learning process after 16 epochs for both V_1 and V_2 .

All classes were tested for both V_1 and V_2 signals. The obtained confusion matrices are presented in Fig. 7 for V_1 and V_2 , while statistical analysis results are shown in Table 5.

Inference on hardware acquired signals confirms that the method is able to generalize even for real-world samples, achieving a macro balanced accuracy of 98.57% for V_1 signals dataset and 100% for V_2 signals dataset. The performance for the two test sets is consistent, demonstrating that the model is able to learn distinctive classification features from both V_1 and V_2 nodes.

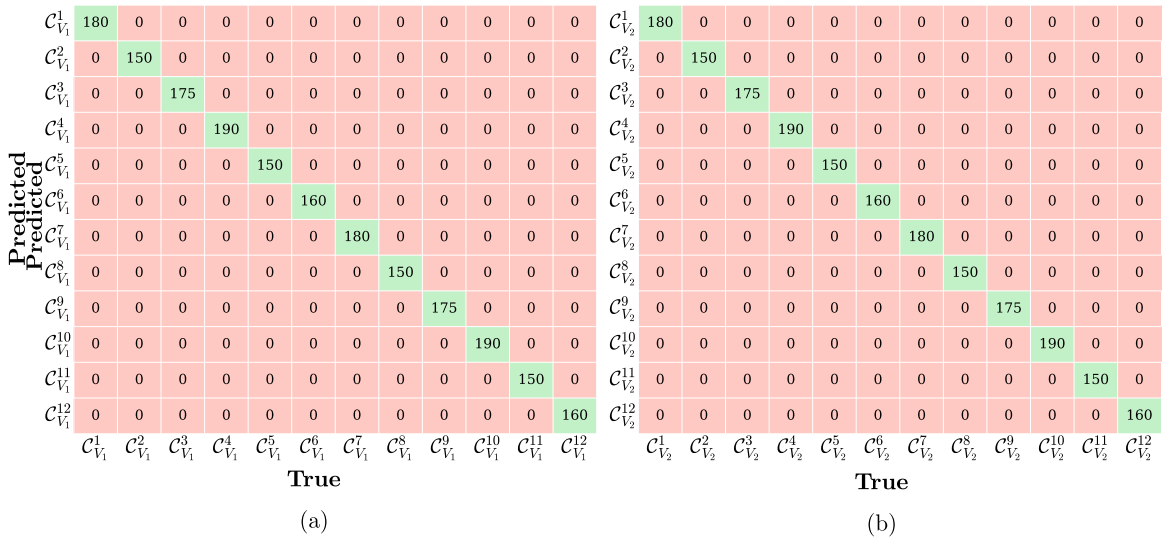


Fig. 8. Confusion matrix obtained on the test set (a) for the V_1 signals and (b) for the V_2 signals of the resistor-swept hardware circuits dataset (Section 6.1).

Table 6

Measured component values and normalized parameters for Circuit 1 and Circuit 2. $\alpha = C_2/C_1$ is computed from the measured values. For $\beta = R_2^2 C_2/L$, the range reflects the minimum and maximum values observed across the six R_2 settings.

	Circuit 1	Circuit 2
L (nH)	19.52	22.20
C_2 (pF)	107.7	106.2
C_1 (pF)	9.72	9.63
R_2 range (Ω)	[1658, 1903]	[1600, 1780]
α	11.08	11.00
β range	[15.16, 19.98]	[12.25, 15.16]

6.1. Validation with resistor-swept hardware circuits

To further expand the experimental validation, a second set of hardware measurements was collected using two independent realizations of Chua’s circuit (hereafter Circuit 1 and Circuit 2). The operating regime is explicitly controlled by sweeping the resistor R_2 across six discrete values, while the remaining passive components (L , C_1 , C_2) are kept physically unchanged. The measured component values and the resulting normalized parameters α and β are reported in Table 6.

The two circuits are treated as independent chaotic sources, each defining six classes corresponding to the six R_2 values, for a total of twelve classes per signal channel, denoted $C^1_{V_s}$ through $C^{12}_{V_s}$ where $s \in \{1, 2\}$ identifies the acquired node (V_1 or V_2), classes $C^1_{V_s}$ through $C^6_{V_s}$ correspond to Circuit 1, and classes $C^7_{V_s}$ through $C^{12}_{V_s}$ to Circuit 2.

From each circuit, voltages V_1 and V_2 at the two capacitor nodes were acquired at a sampling frequency of $f_s = 400$ kHz. Signal preprocessing followed the same pipeline described in Section 6. For each signal channel, the twelve classes were pooled into a single dataset and split into training (60%), validation (20%), and test (20%) subsets using stratified random sampling, consistent with the procedure of Section 3. The total dataset comprises approximately 5000 samples per channel, with a slight class imbalance across cases (ranging from 750 to 950 samples per class). A single YOLOv11-Medium classifier per channel was trained, with same option of Section 3, on the combined twelve-class dataset.

Classification results on the test sets are reported in Fig. 8 and Table 7. Both classifiers, for V_1 and V_2 , achieve perfect separation of all twelve classes, with macro balanced accuracy of 100% and fully diagonal confusion matrices with zero misclassifications.

7. Discussion

7.1. Classification performance on simulated datasets

This work shows that discriminating multiple operating regimes of the same chaotic generator under a fixed topology is not only feasible but can achieve high accuracy when chaotic time series are rendered in the time-frequency domain. Across the considered

Table 7

Classification performance metrics for the resistor-swept hardware circuits dataset. Classes $C_{V_1}^1 - C_{V_1}^6$ correspond to Circuit 1 and $C_{V_2}^7 - C_{V_2}^{12}$ to Circuit 2. All results are reported on the test set.

Class	Samples	Precision	Recall	Balanced accuracy
$C_{V_1}^1$	180	100%	100%	100%
$C_{V_1}^2$	150	100%	100%	100%
$C_{V_1}^3$	175	100%	100%	100%
$C_{V_1}^4$	190	100%	100%	100%
$C_{V_1}^5$	150	100%	100%	100%
$C_{V_1}^6$	160	100%	100%	100%
$C_{V_2}^7$	180	100%	100%	100%
$C_{V_2}^8$	150	100%	100%	100%
$C_{V_2}^9$	175	100%	100%	100%
$C_{V_2}^{10}$	190	100%	100%	100%
$C_{V_2}^{11}$	150	100%	100%	100%
$C_{V_2}^{12}$	160	100%	100%	100%
$C_{V_2}^1$	180	100%	100%	100%
$C_{V_2}^2$	150	100%	100%	100%
$C_{V_2}^3$	175	100%	100%	100%
$C_{V_2}^4$	190	100%	100%	100%
$C_{V_2}^5$	150	100%	100%	100%
$C_{V_2}^6$	160	100%	100%	100%
$C_{V_2}^7$	180	100%	100%	100%
$C_{V_2}^8$	150	100%	100%	100%
$C_{V_2}^9$	175	100%	100%	100%
$C_{V_2}^{10}$	190	100%	100%	100%
$C_{V_2}^{11}$	150	100%	100%	100%
$C_{V_2}^{12}$	160	100%	100%	100%

datasets, the trained YOLOv11-based classifiers learn discriminative representations from spectrograms derived from both circuit nodes (x and y), with closely matched performance between channels. This consistency suggests that the models are not exploiting node-specific artifacts, but rather capture regime-dependent time-frequency structure that is observable at both voltages.

In the first Case Study, the classes are designed to separate parameter-driven regimes under controlled conditions. Classes $S_{\{x|y\}}^1$ to $S_{\{x|y\}}^3$ correspond to three fixed sets of nonlinear coefficients, while $S_{\{x|y\}}^4$ intentionally introduces higher intra-class variability by sampling nonlinear parameters within broad chaos-preserving ranges overlapping the other regimes. Despite this deliberate “nearby-regime” construction, the confusion matrices remain strongly diagonal for both x and y , indicating that the classifier rarely confuses the variable class with the fixed ones. This is a non-trivial result because $S_{\{x|y\}}^4$ is explicitly designed to reduce separability and to test whether the network extracts regime-relevant cues rather than memorizing fixed templates.

In the second Case Study, the parameter space is widened and densely sampled to generate 200 distinct chaotic classes per channel, substantially increasing the number of closely related regimes and thus providing a harsher discrimination setting. Even in this large-scale regime sampling, macro-level scores remain high, and the per-class curves of balanced accuracy, precision, and recall stay consistently elevated across most classes, with only a limited subset exhibiting noticeable dips. Importantly, the per-class trends are highly similar for x and y , reinforcing the interpretation that both channels encode comparable discriminative information from the perspective of the model.

Additional targeted checks further support the intended semantics of the class design. When digitally summing signals from different fixed-parameter regimes in the first case study, the resulting hybrid trajectories are consistently assigned to $S_{\{x|y\}}^4$, the heterogeneous “none-of-the-previous” class. This behavior is aligned with the role of $S_{\{x|y\}}^4$ as a distribution-shift class and suggests that the classifier reacts coherently to mixed spectral content rather than forcing hybrids into the closest fixed-template class.

7.2. Time-frequency explainability and signature validation

Beyond classification accuracy, the analysis in Section 5.1 provides evidence that the pipeline is learning structured time-frequency features of the signatures. Grad-CAM localization (computed on the second Case Study for statistical representativeness) highlights salient spectrogram regions that are not randomly scattered, and PERMANOVA rejects the null hypothesis of no between-class differences in the Grad-CAM-derived time-frequency distributions. Together, these results support the presence of regime-dependent saliency patterns that are stable enough to be detected at scale, and associated with the notion of “chaotic signatures” as discriminative time-frequency structures linked to parameter configurations.

7.3. Hardware validation and proof of concept transferability

Validation on hardware-acquired Chua's circuit signals indicates that the learned representations are not tied exclusively to simulation artifacts. Despite additional variability due to component tolerances, measurement noise and other non-idealities, the classifier maintains strong performance and comparable behavior across V_1 and V_2 in both experimental datasets presented in Sections 6 and 6.1. While the experimental datasets remain smaller than the simulated ones, these results provide a proof of concept that the method is transferable to physical circuits. Future experimental work will include a systematic evaluation of temperature effects and filtering effects on the learned signatures, alongside a broader validation campaign across additional circuit realizations and operating conditions.

7.4. YOLOv11 architectural choice and motivation

The choice of YOLOv11 for the classification of chaotic signal spectrograms is motivated by its demonstrated effectiveness in spectrogram-based analysis across several signal processing domains, including RF signal identification, distributed fiber sensing, and audio analysis [17–19,43]. This track record suggests that its backbone is well-suited to extracting discriminative features from complex time-frequency representations, a property directly relevant to the classification task addressed in this work, as confirmed by the high classification accuracy achieved across all case studies presented in this paper. Furthermore, the choice of a backbone with strong multi-scale feature extraction capabilities [44] was deliberate, as capturing hierarchical time-frequency representations improves the spatial localization of Grad-CAM saliency maps, enabling a more precise identification of the discriminative spectral regions associated with each chaotic signature.

To further contextualize this architectural choice, we trained ResNet50 and EfficientNet-B0 on the same STFT spectrograms under identical training conditions on a single NVIDIA GH200 480 GB GPU (CUDA 12.6). All three architectures achieved comparable macro-averaged balanced accuracy on the test set of the simple case study (Section 3), averaged over the two models trained on x -derived and y -derived spectrograms (ResNet50: 99.56%, EfficientNet-B0: 99.14%, YOLOv11: 99.30%), with differences below 0.5% that are practically negligible. In terms of computational cost, all three architectures showed comparable peak memory usage (approximately 9.5–10.1 GB on the same hardware). The most significant difference emerged in training convergence: YOLOv11 required only 6.7 min averaged over the two training runs (one for x -derived spectrograms and one for y -derived spectrograms), compared to 9.1 and 13.2 min for EfficientNet-B0 and ResNet50 respectively. This advantage becomes practically relevant when scaling to the 200-class setting of the more complex case study (Section 4), where the reduced training time represents a meaningful gain in the overall experimental workflow.

7.5. Extension to different hardware configurations

An additional promising direction concerns the extension of the proposed framework to other chaotic circuit families and hardware implementations. This includes both different chaotic oscillator topologies and variants of Chua's circuit itself, such as memristive implementations in which the classical nonlinear resistor is replaced by or coupled with a memristor element [45]. Such variants exhibit richer attractor morphologies and an expanded parameter-dependent dynamical repertoire compared to the classical formulation, potentially offering a broader and more complex signature space for fingerprinting applications. Investigating whether the STFT-based pipeline can extract stable, discriminative signatures across different circuit families and hardware configurations represents a natural and theoretically motivated extension of the present work, further connecting it to the broader landscape of chaotic hardware security applications [46].

8. Conclusion

We presented a time-frequency deep-learning pipeline to classify chaotic regimes within a single circuit family (Chua's circuit) by treating parameterizations as labeled regimes and learning discriminative patterns directly from STFT spectrograms. Across two simulated case studies (from controlled fixed-parameter regimes to 200-class randomized regimes) and two experimental hardware datasets, the proposed approach achieved consistently strong performance and closely matched results between both x and y signals.

Explainability and statistics provide additional support for the "signature" hypothesis. On the large-scale randomized dataset, Grad-CAM highlights localized time-frequency regions that contribute most to classification decisions, and PERMANOVA confirms that the corresponding saliency distributions differ significantly across classes. Overall, these findings suggest that regime-specific time-frequency patterns can serve as candidate features of different regime's signature. Future work will expand experimental validation across additional devices and circuit configurations. We will also compare Grad-CAM with complementary explainability methods to assess the robustness of the signature features that most consistently support regime discrimination, and investigate whether the salient time-frequency regions exhibit consistent patterns across different operating conditions.

CRediT authorship contribution statement

Sara Becchi: Writing – review & editing, Writing – original draft, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Elisabetta Spinazzola:** Writing – review & editing, Writing – original draft, Visualization. **Serhii Haliuk:** Writing – review & editing, Resources, Data curation. **Fernando Corinto:** Writing – review & editing, Supervision. **Leon O. Chua:** Supervision. **Fabio Pareschi:** Writing – review & editing, Visualization, Validation, Supervision, Project administration, Funding acquisition, Conceptualization. **Dmytro Vovchuk:** Writing – review & editing, Supervision, Resources, Data curation. **Jacopo Secco:** Writing – review & editing, Validation, Supervision, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Sara Becchi reports financial support was provided by Polytechnic of Turin. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU. This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

Data availability

Data will be made available on request.

References

- [1] Giakoumis A, Volos CK, Munoz-Pacheco JM, del Carmen Gomez-Pavon L, Stouboulos IN, Kyprianidis IM. Text encryption device based on a chaotic random bit generator. In: 2018 IEEE 9th latin American symposium on circuits & systems. LASCAS, IEEE; 2018, p. 1–5. <http://dx.doi.org/10.1109/LASCAS.2018.8399974>.
- [2] Alvarez G, Li S. Some basic cryptographic requirements for chaos-based cryptosystems. *Int J Bifurc Chaos* 2006;16(08):2129–51. <http://dx.doi.org/10.1142/S0218127406015970>.
- [3] Pareschi F, Setti G, Rovatti R. Implementation and testing of high-speed CMOS true random number generators based on chaotic systems. *IEEE Trans Circuits Syst I Regul Pap* 2010;57(12):3124–37. <http://dx.doi.org/10.1109/TCSI.2010.2052515>.
- [4] Merah L, Ali-Pacha A, Said NH, Mamat M. A pseudo random number generator based on the chaotic system of chua's circuit, and its real time FPGA implementation. *Appl Math Sci* 2013;7(55):2719–34. [10.12988/ams.2013.13242](https://doi.org/10.12988/ams.2013.13242).
- [5] Chua LO. The genesis of Chua's circuit. Electronics Research Laboratory, College of Engineering, University of California; 1992.
- [6] Chen L. A framework to enhance security of physically unclonable functions using chaotic circuits. *Phys Lett A* 2018;382(18):1195–201. <http://dx.doi.org/10.1016/j.physleta.2018.03.012>.
- [7] Soltanieh N, Norouzi Y, Yang Y, Karmakar NC. A review of radio frequency fingerprinting techniques. *IEEE J Radio Freq Identif* 2020;4(3):222–33. <http://dx.doi.org/10.1109/JRFID.2020.2968369>.
- [8] Gomes I, Korneta W, Stavrinides SG, Picos R, Chua LO. Experimental observation of chaotic hysteresis in chua's circuit driven by slow voltage forcing. *Chaos Solitons Fractals* 2023;166:112927. <http://dx.doi.org/10.1016/j.chaos.2022.112927>.
- [9] Haliuk S, Vovchuk D, Spinazzola E, Secco J, Brobrovs V, Corinto F. A deterministic chaos-model-based Gaussian noise generator. *Electronics* 2024;13(7):1387. <http://dx.doi.org/10.20944/preprints202403.0384.v1>.
- [10] Matsumoto T. A chaotic attractor from chua's circuit. *IEEE Trans Circuits Syst* 2003;31(12):1055–8. <http://dx.doi.org/10.1109/TCS.1984.1085459>.
- [11] Jia B, Wu H, Guo K. Chaos theory meets deep learning: A new approach to time series forecasting. *Expert Syst Appl* 2024;255:124533. <http://dx.doi.org/10.1016/j.eswa.2024.124533>.
- [12] Akmeşe ÖF, Emin B, Alaca Y, Karaca Y, Akgül A. The time series classification of discrete-time chaotic systems using deep learning approaches. *Mathematics* 2024;12(19):3052. <http://dx.doi.org/10.3390/math12193052>.
- [13] Giuseppi A, Menegatti D, Pietrabissa A. Identifying chaotic dynamics in noisy time series through multimodal deep neural networks. *Mach Learn: Sci Technol* 2024;5(3):035059. <http://dx.doi.org/10.1088/2632-2153/ad7190>.
- [14] Mukhopadhyay S, Banerjee S. Learning dynamical systems in noise using convolutional neural networks. *Chaos: An Interdiscip J Nonlinear Sci* 2020;30(10). <http://dx.doi.org/10.1063/5.0009326>.
- [15] Uzun S, Kaçar S, Arıcıoğlu B. Deep learning based classification of time series of chaotic systems over graphic images. *Multimedia Tools Appl* 2024;83(3):8413–37. <http://dx.doi.org/10.1007/s11042-023-15944-3>.
- [16] Arıcıoğlu B, Uzun S, Kaçar S. Deep learning based classification of time series of chen and rössler chaotic systems over their graphic images. *Phys D: Nonlinear Phenom* 2022;435:133306. <http://dx.doi.org/10.1016/j.physd.2022.133306>.
- [17] Delleji T, Slimeni F. RF-YOLO: a modified YOLO model for UAV detection and classification using RF spectrogram images. *Telecommun Syst* 2025;88(1):33. <http://dx.doi.org/10.1007/s11235-025-01264-4>.
- [18] Zbezhkhovska U. Deepfake audio detection using yolov8 with mel-spectrogram analysis: a cross-dataset evaluation. *Radio Electron Comput Sci Control* 2025;153–63. <http://dx.doi.org/10.15588/1607-3274-2025-1-14>.
- [19] Huynh-The T, Huynh P-T, Pham Q-V. Spec-YOLO: An efficient deep network for spectrogram-based signals identification. *IEEE Wirel Commun Lett* 2025. <http://dx.doi.org/10.1109/LWC.2025.3598052>.
- [20] Matsumoto T, Chua L, Komuro M. The double scroll. *IEEE Trans Circuits Syst* 2003;32(8):797–818. <http://dx.doi.org/10.1109/TCS.1986.1085869>.
- [21] Kuznetsov N, Mokaev T, Ponomarenko V, Seleznev E, Stankevich N, Chua L. Hidden attractors in chua circuit: mathematical theory meets physical experiments. *Nonlinear Dynam* 2023;111(6):5859–87. <http://dx.doi.org/10.1007/s11071-022-08078-y>.
- [22] Kennedy MP. Three steps to chaos. II. A chua's circuit primer. *IEEE Trans Circuits Syst I* 2002;40(10):657–74. <http://dx.doi.org/10.1109/81.246141>.
- [23] Dormand JR, Prince PJ. A family of embedded runge-kutta formulae. *J Comput Appl Math* 1980;6(1):19–26. [http://dx.doi.org/10.1016/0771-050X\(80\)90013-3](http://dx.doi.org/10.1016/0771-050X(80)90013-3).
- [24] Benettin G, Galgani L, Giorgilli A, Strelcyn J-M. Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part 1: Theory. *Meccanica* 1980;15(1):9–20. <http://dx.doi.org/10.1007/BF02128236>.
- [25] Geist K, Parlitz U, Lauterborn W. Comparison of different methods for computing Lyapunov exponents. *Progr Theoret Phys* 1990;83(5):875–93. <http://dx.doi.org/10.1143/PTP.83.875>.
- [26] Sprott J. Some simple chaotic jerk functions. *Am J Phys* 1997;65(6):537–43. <http://dx.doi.org/10.1119/1.18585>.
- [27] Sano M, Sawada Y. Measurement of the Lyapunov spectrum from a chaotic time series. *Phys Rev Lett* 1985;55(10):1082. <http://dx.doi.org/10.1103/PhysRevLett.55.1082>.

- [28] Wolf A, Swift JB, Swinney HL, Vastano JA. Determining Lyapunov exponents from a time series. *Phys D: Nonlinear Phenom* 1985;16(3):285–317. [http://dx.doi.org/10.1016/0167-2789\(85\)90011-9](http://dx.doi.org/10.1016/0167-2789(85)90011-9).
- [29] Rosenstein MT, Collins JJ, De Luca CJ. A practical method for calculating largest Lyapunov exponents from small data sets. *Phys D: Nonlinear Phenom* 1993;65(1–2):117–34. [http://dx.doi.org/10.1016/0167-2789\(93\)90009-P](http://dx.doi.org/10.1016/0167-2789(93)90009-P).
- [30] Abarbanel HD, Brown R, Sidorowich JJ, Tsimring LS. The analysis of observed chaotic data in physical systems. *Rev Modern Phys* 1993;65(4):1331. <http://dx.doi.org/10.1103/RevModPhys.65.1331>.
- [31] Cohen L. Time-frequency distributions—a review. *Proc IEEE* 2002;77(7):941–81. <http://dx.doi.org/10.1109/5.30749>.
- [32] Hammond JK, White PR. The analysis of non-stationary signals using time-frequency methods. *J Sound Vib* 1996;190(3):419–47. <http://dx.doi.org/10.1006/jsvi.1996.0072>.
- [33] Allen JB, Rabiner LR. A unified approach to short-time Fourier analysis and synthesis. *Proc IEEE* 2005;65(11):1558–64. <http://dx.doi.org/10.1109/PROC.1977.10770>.
- [34] Harris FJ. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc IEEE* 2005;66(1):51–83. <http://dx.doi.org/10.1109/PROC.1978.10837>.
- [35] Khanam R, Hussain M. Yolov11: An overview of the key architectural enhancements. 2024, <http://dx.doi.org/10.48550/arXiv.2410.17725>, arXiv preprint arXiv:2410.17725.
- [36] Zhuang F, Qi Z, Duan K, Xi D, Zhu Y, Zhu H, Xiong H, He Q. A comprehensive survey on transfer learning. *Proc IEEE* 2020;109(1):43–76. <http://dx.doi.org/10.1109/TMI.2016.2528162>.
- [37] Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE international conference on computer vision*. 2017, p. 618–26. <http://dx.doi.org/10.48550/arXiv.1610.02391>.
- [38] Faruqui Z, McIntire MS, Dubey R, McEntee J. Explainability of CNN based classification models for acoustic signal. 2025, <http://dx.doi.org/10.48550/arXiv.2509.08717>, arXiv preprint arXiv:2509.08717.
- [39] Bajwa MKZ, Castiglione A, Pero C. Mel spectrogram-based CNN framework for explainable audio deepfake detection. In: *International conference on advanced information networking and applications*. Springer; 2025, p. 407–16. http://dx.doi.org/10.1007/978-3-031-87784-1_37.
- [40] O’Shea TJ, Roy T, Erpek T. Spectral detection and localization of radio events with learned convolutional neural features. In: *2017 25th European signal processing conference. EUSIPCO, IEEE*; 2017, p. 331–5. <http://dx.doi.org/10.23919/EUSIPCO.2017.8081223>.
- [41] Anderson MJ. A new method for non-parametric multivariate analysis of variance. *Austral Ecol* 2001;26(1):32–46. <http://dx.doi.org/10.1111/j.1442-9993.2001.01070.pp.x>.
- [42] Huttenlocher DP, Klanderman GA, Rucklidge WJ. Comparing images using the hausdorff distance. *IEEE Trans Pattern Anal Mach Intell* 1993;15(9):850–63. <http://dx.doi.org/10.1109/34.232073>.
- [43] Xu W, Yu F, Liu S, Xiao D, Hu J, Zhao F, Lin W, Wang G, Shen X, Wang W, et al. Real-time multi-class disturbance detection for Φ -OTDR based on YOLO algorithm. *Sensors* 2022;22(5):1994. <http://dx.doi.org/10.3390/s22051994>.
- [44] He L-h, Zhou Y-z, Liu L, Cao W, Ma J-h. Research on object detection and recognition in remote sensing images based on YOLOv11. *Sci Rep* 2025;15(1):14032. <http://dx.doi.org/10.1038/s41598-025-96314-x>.
- [45] Chua L. Memristor—the missing circuit element. *IEEE Trans Circuit Theory* 1971;18(5):507–19. <http://dx.doi.org/10.1109/TCT.1971.1083337>.
- [46] Lv S, Liu J, Geng Z. Application of memristors in hardware security: A current state-of-the-art technology. *Adv Intell Syst* 2021;3(1):2000127. <http://dx.doi.org/10.1002/aisy.202000127>.