PURPOSE-LED
PUBLISHING™

**PAPER • OPEN ACCESS**

# Identifying chaotic dynamics in noisy time series through multimodal deep neural networks

To cite this article: Alessandro Giuseppi *et al* 2024 *Mach. Learn.: Sci. Technol.* **5** 035059

View the article online for updates and enhancements.

## MACHINE LEARNING
Science and Technology

**PAPER**

# Identifying chaotic dynamics in noisy time series through multimodal deep neural networks

Alessandro Giuseppi[*] , Danilo Menegatti and Antonio Pietrabissa

Department of Computer, Control and Management Engineering of the University of Rome La Sapienza, Via Ariosto 25, 00185 Rome, Italy

[*] Author to whom any correspondence should be addressed.

**E-mail:** giuseppi@diag.uniroma1.it

## Abstract

Chaos detection is the problem of identifying whether a series of measurements is being sampled from an underlying set of chaotic dynamics. The unavoidable presence of measurement noise significantly affects the performance of chaos detectors, as discerning chaotic dynamics from stochastic signals becomes more challenging. This paper presents a computationally efficient multimodal deep neural network tailored for chaos detection by combining information coming from the analysis of time series, recurrence plots and spectrograms. The proposed approach is the first one suitable for multi-class classification of chaotic systems while being robust with respect to measurement noise, and is validated on a dataset of 15 different chaotic and non-chaotic dynamics subject to white, pink or brown colored noise.

## 1. Introduction

Chaos is a behavior that affects several deterministic nonlinear dynamical systems which exhibit complex, random-looking, evolutions due to a significant sensitivity to small changes in their initial conditions.

When studying a time series of observations, the problem of *chaos detection* consists in understanding whether the underlying dynamics is of a chaotic nature or not, and it is crucial for the correct analysis of the system and to predict its future evolution. In fact, discerning deterministic, albeit chaotic, dynamics from stochastic processes opens up several opportunities in fields that can benefit from having a reliable prediction regarding the future evolution of a complex system, such as is the case for financial market analysis, weather forecast, epidemiology and neurosciences [1] with significant impacts on entire industrial fields and markets. Detecting chaos in empirical data is a particularly complex task, as the unavoidable presence of noise affecting the observations may render the discernment of chaotic dynamics from periodic deterministic trajectories and stochastic signals particularly challenging [2].

In this direction, this paper presents an ad-hoc multimodal deep neural network (DNN) architecture tailored for detecting and classifying chaotic dynamics in noisy time series. The proposed DNN employs standard tools for chaos analysis, namely recurrence plots and spectrograms, to recognize the underlying dynamics from partial state observations corrupted by stochastic noise. The proposed DNN architecture is then validated on a dataset derived from the one originally presented in [3] to discern among 15 different dynamics subject to various levels and types of stochastic noise.

This study improves existing literature in several aspects:

- The proposed DNN treats chaos detection as a multi-class classification problem, whereas most of the machine learning-based approaches in the literature only focus on a single set of chaotic dynamics following a binary approach.
- The developed system shows a high level of tolerance against measurement noise.
- The DNN allows for computationally efficient chaos detection, as the computing burden is limited to its off-line training phase.

The remainder of the paper is organized as follows: section 2 presents a survey of the relevant literature; section 3 introduces some needed preliminaries and definitions; section 4 describes and explains the proposed neural network architecture; section 5 discusses the results of the validation of the designed DNN carried out on a recent public dataset; section 6 draws the conclusions and highlights possible future works.

*Notation*: $\mathbf{A}_{i,j}$ denotes the *j*th element of the *i*th row of the matrix $\mathbf{A}$; the operator $||\cdot||$ denotes the $L^2$-norm.

## 2. Related works

Chaos theory is the branch of systems theory that deals with the analysis and study of chaotic dynamical systems and their trajectories. A chaotic system is characterized by deterministic dynamics that amplify exponentially small perturbations of its initial conditions, meaning that small differences in the system's initial state lead to significantly different trajectories causing what is commonly known as the *butterfly effect* [2].

Over the years, due to the ubiquitous presence of chaos in natural processes, chaos theory has caused a significant impact in almost all fields of engineering, spacing from fluid and plasma dynamics [4] to neuroscience [5], optics [6] and medicine [7]. Chaos detection is the problem of understanding whether or not a set of observations is sampled from a dynamical system that is characterized by chaotic deterministic dynamics, allowing for proper system analysis and potentially the design of adequate control laws.

Assuming *a priori* knowledge on the structure of the system dynamics, a first approach for chaos detection involves the determination of which regions of the system's parameter space lead to chaotic behaviors. In this setting, a common approach to determine such regions is based on the analytical determination of the maximal Lyapunov characteristic exponent (LCE) [8] to estimate how fastly neighboring trajectories diverge with time.

A different approach to chaos detection can be followed in a data-driven setting when no information is available on the system dynamics, and one is interested in determining whether a given time series of observations has been produced by a chaotic system. In such a scenario, it is still possible to estimate the LCE directly from the time series [8, 9], but the estimation sensitivity to measurement noise and its limitations in terms of sample and computational requirements has led scientists to design more robust and efficient solutions [2, 10]. In this direction, over the last few years several works investigated how machine and deep learning solutions [3, 11–16] may contribute to the chaos detection problem. The authors in [11–13] tested various DNN architectures as binary classifiers fed directly with the observed time series to detect various different chaotic dynamics. In [11, 12] it was assumed to have *a priori* knowledge on the system dynamics, as each DNN was specialized for a single system, while [13] tests how a DNN trained for binary chaos detection of the Chirikov standard map performs in recognizing chaos also on the Logistic map and Lorenz system. In [3] a different approach is followed and both convolutional and recurrent neural networks are tested to classify a time series as chaotic or non-chaotic. Works such as [16–18] pursue a LCE-based analysis by combining DNNs with the typical methods for chaos detection: in [17, 18] an empirical data-driven solution was developed to provide chaos detection without any information on the system dynamics; in [16] 2D-bifurcation diagrams were reconstructed with various DNNs to identify chaos over the parameter space of a given system.

Following a successful chaos detection, a typical case study involves the prediction of a chaotic system time evolution, as it may be used to plan/optimize the applied control or the system structure itself. We mention that DNN have found significant applications also in this direction [19–25], as they offer computationally-efficient tools with demonstrated high accuracy.

As mentioned, one of the main limitations of most of the chaos detection approaches available in the literature is due to the natural presence of additive noise in empirical measurements. To limit the impact of noise on chaos detection, works such as [14, 15] focused on discerning purely stochastic time series from trajectories of chaotic systems affected by noise. In [14], several DNN are tested for chaos detection for a total of six chaotic systems, but interestingly the authors trained all DNN as binary classifiers, hence assuming *a priori* knowledge on the system dynamics. We mention that the authors of [14], similarly to [13], also test whether their better-performing DNN is capable of discerning different chaotic dynamics with respect to the one it was specifically trained for, but observed limited accuracy that highlights an overall low generalization. A different approach is followed by the authors of [15], which consider four different chaotic systems corrupted by three types of additive colored noise (white, pink or brown). Chaos detection is then solved by designing a DNN to classify either the time series as chaotic or stochastic, identifying also the noise affecting the time series. Similarly to the present paper, in [15] the DNN is fed with recurrence plots and spectrograms, but differently from our approach, the authors discard the time series and employ a single DNN channel and provide no further classification regarding the specific chaotic dynamics that generated the data.

The present paper proposes a specialized DNN capable of first analyzing in parallel, and then combining information coming from recurrence plots, spectrograms and time series to solve chaos detection in the presence of significant levels of measurement noise. The DNN has been designed to recognize which specific system dynamics is the most likely to have generated the sampled data among the five chaotic and ten non-chaotic systems presented in [3, 26].

The proposed method is the first one capable of multi-class classification of chaotic systems while also demonstrating robustness against measurement noise, as shown by the simulations of section 5.

# 3. Preliminaries

This section introduces and provides the definitions of the main tools employed for the analysis of the time series and the design of the proposed DNN.

### 3.1. Problem description: chaos detection in noisy time series

For the design of DNNs we will adhere to the following assumptions: (i) the time series available for our analysis will be limited to a single state component, as in the example in figures 1 and 2; (ii) the system dynamics are sampled from one out of 15 different sets of dynamics.

The former assumption allows us to feed our system with time series sampled from a broad class of systems, independently from their state dimension. We mention that extending the proposed architecture to the multi-dimensional case is almost seamless, as it would be sufficient to either stack equivalent inputs (e.g. all recurrence plots) before feeding them to the convolutional layers or to add some additional channels in the DNN. Also, the number of dynamics is arbitrary and is derived from the dataset we employed, [26], which is significantly higher than the typical number of dynamics considered in similar studies such as [12, 15].

### 3.2. Chaotic dynamics

The five chaotic dynamics considered for our design are derived from [3] and detailed in the following:

(i)  Ueda oscillator

$$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = -x_1^3 - bx_2 + A\sin(\Omega t), \end{cases} \tag{1}$$

with $\Omega = 1, A = 7.5, b = 0.05$.

(ii)  Lorenz attractor

$$\begin{cases} \dot{x}_1 = -\sigma x_1 + \sigma x_2 \\ \dot{x}_2 = \rho x_1 - x_2 - x_1 x_3 \\ \dot{x}_3 = -\beta x_3 + x_1 x_2 \end{cases} \tag{2}$$

with $\sigma = 10, \beta = 8/3, \rho = 28$.

(iii)  Rossler attractor

$$\begin{cases} \dot{x}_1 = -x_2 - x_3 \\ \dot{x}_2 = x_1 + ax_2 \\ \dot{x}_3 = b + x_3(x_1 - c) \end{cases} \tag{3}$$

with $a = 0.2, b = 0.2, c = 5.7$.
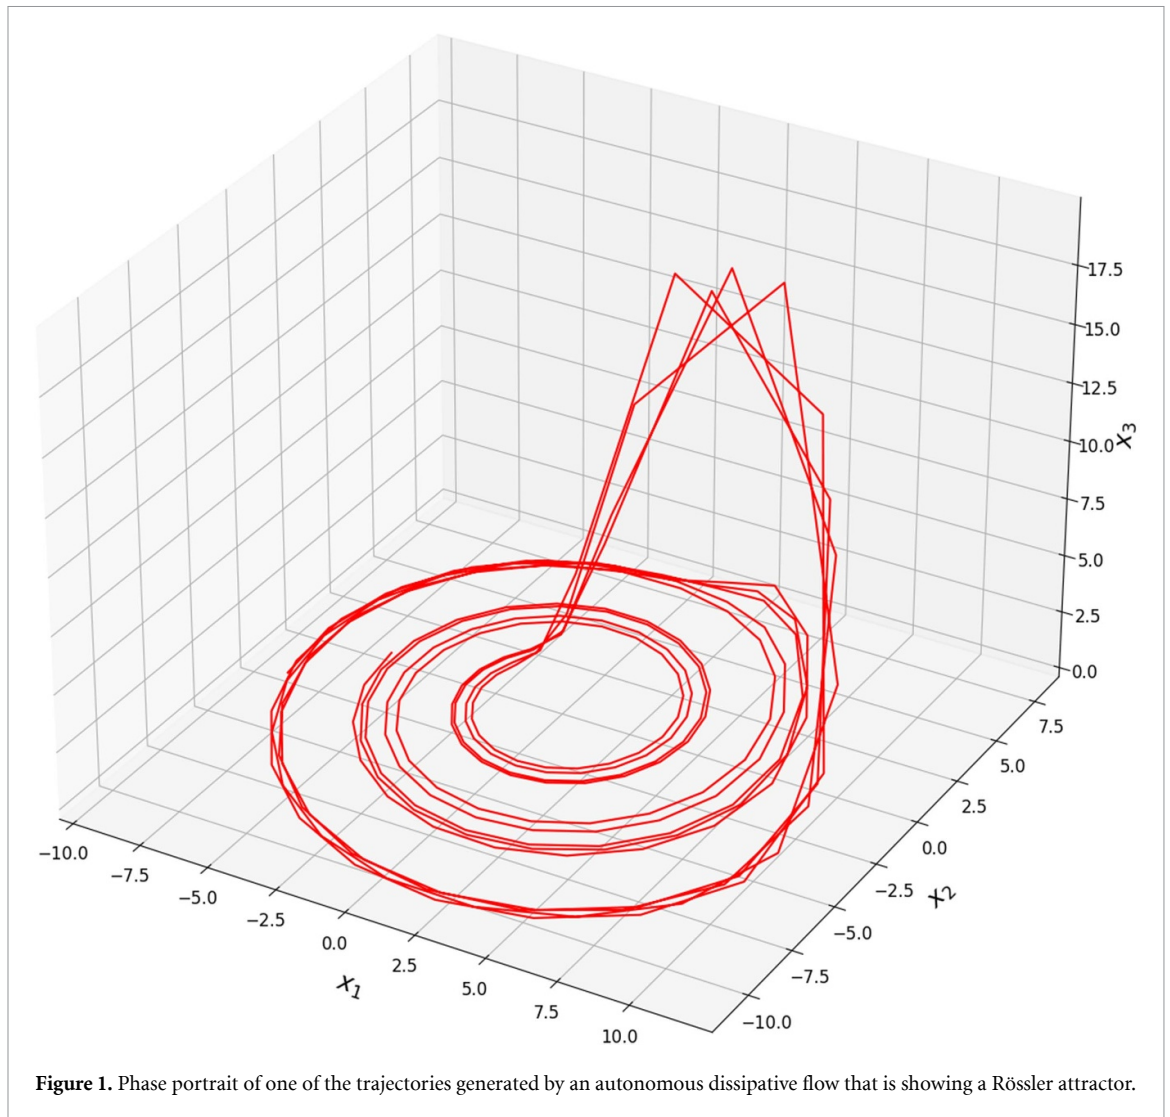
(iv)  Halvorsen attractor

$$\begin{cases} \dot{x}_1 = -ax_1 - 4x_2 - 4x_3 - x_2{}^2 \\ \dot{x}_2 = -4x_1 - ax_2 - 4x_3 - x_3{}^2 \\ \dot{x}_3 = -4x_1 - 4x_2 - ax_3 - x_1{}^2 \end{cases} \tag{4}$$

with $a = 1.27$.

(v)  Rucklidge attractor

$$\begin{cases} \dot{x}_1 = -kx_1 + \lambda x_2 - x_2 x_3 \\ \dot{x}_2 = x_1 \\ \dot{x}_3 = -x_3 + x_2^2 \end{cases} \tag{5}$$

with $k = 2, \lambda = 6.7$.

**Figure 1.** Phase portrait of one of the trajectories generated by an autonomous dissipative flow that is showing a Rössler attractor.

Regarding the non-chaotic dynamics considered, [3] includes a broad range of dynamical systems with periodic, quasi-periodic and non-periodic behaviors.

### 3.3. Recurrence plots

Recurrence plots [27] were originally presented in the 80's to provide a clear visualization of how a system trajectory evolves in its phase space in terms of how often it *recurs* in a neighborhood of a state it previously visited. Let $x(t)$ be the state of a dynamical system at time $t$; a trajectory of state measures $x(1), x(2), \ldots, x(N)$ is used to define the so-called recurrence matrix $\mathbf{R}$, whose elements are defined as [27]:

$$\mathbf{R}_{i,j} = \begin{cases} 1 & \text{if } x(i) \approx x(j) \\ 0 & \text{otherwise} \end{cases}, \tag{6}$$

where the operator $\approx$ is used to represents equality up to a threshold $\epsilon > 0$ in the sense that $||x(i) - x(j)|| < \epsilon$.

A recurrence plot is then a way to visualize graphically, encoding each pixel in black or white, the entries of the matrix (6) and can be used to identify periodic, quasi-periodic and recurrence towards attractor behaviors, depending on the particular structures that appear (e.g. diagonal lines, isolated points, patterns).

Several variations of recurrence plots have been proposed to capture different characteristics of the system dynamics by considering different norms, definitions of the distance among two states $x(i), x(j)$ and/or thresholds [27], impacting on the very definition of the concept of recurrence used in their analysis. This study employs the so-called unthresholded recurrence plots [28] (e.g. figure 3), also known as distance plots, which enriches the information visualized by (6) by removing the threshold $\epsilon$ and instead encoding in a range of colors the distance $||x(i) - x(j)||$, i.e.:

$$\mathbf{R}_{i,j} = ||x(i) - x(j)||. \tag{7}$$

**Figure 2.** (top) Time series of measurements relative to the $x_1$ state for the example of figure 1; (bottom) same time series affected by a white Gaussian noise with a signal-to-noise ratio of 5 dB.

The rationale behind our choice is to provide the DNN with as much information as possible, as the threshold introduced in (6) discards any insight on the distance magnitude and relies on the arbitrarily set $\epsilon$. In fact, providing the DNN with the unthresholded recurrence plot allows it to identify autonomously the most effective features to consider to define recurrence.

From figure 3 we see how noise affects recurrence plots by apparently blurring them with random vertical and horizontal lines. However, we note how the general location of peaks and valleys is preserved, suggesting that information about recurrence is preserved.

### 3.4. Spectrograms

Spectrograms are among the most broadly used visualization tools to analyze time series as they show how the spectrum of frequencies of the given signal varies with time.

The typical way of representing a spectrogram is as a heatmap $\mathbf{S}$, as the ones depicted in figure 4, in which for a given time $j$ the elements/pixels $\mathbf{S}_{i,j}$ encode the power spectral density (in dB) for the various frequencies $i$.

From figure 4 it is clear that noise re-distributes power over the various frequencies. Having considered three different colored noises (white, pink and brown), the effect they cause on the original signal spectrum follows their power distribution, which is uniform for white noise, inversely proportional to frequency for pink noise and inversely proportional to the square of the frequency for brown noise.
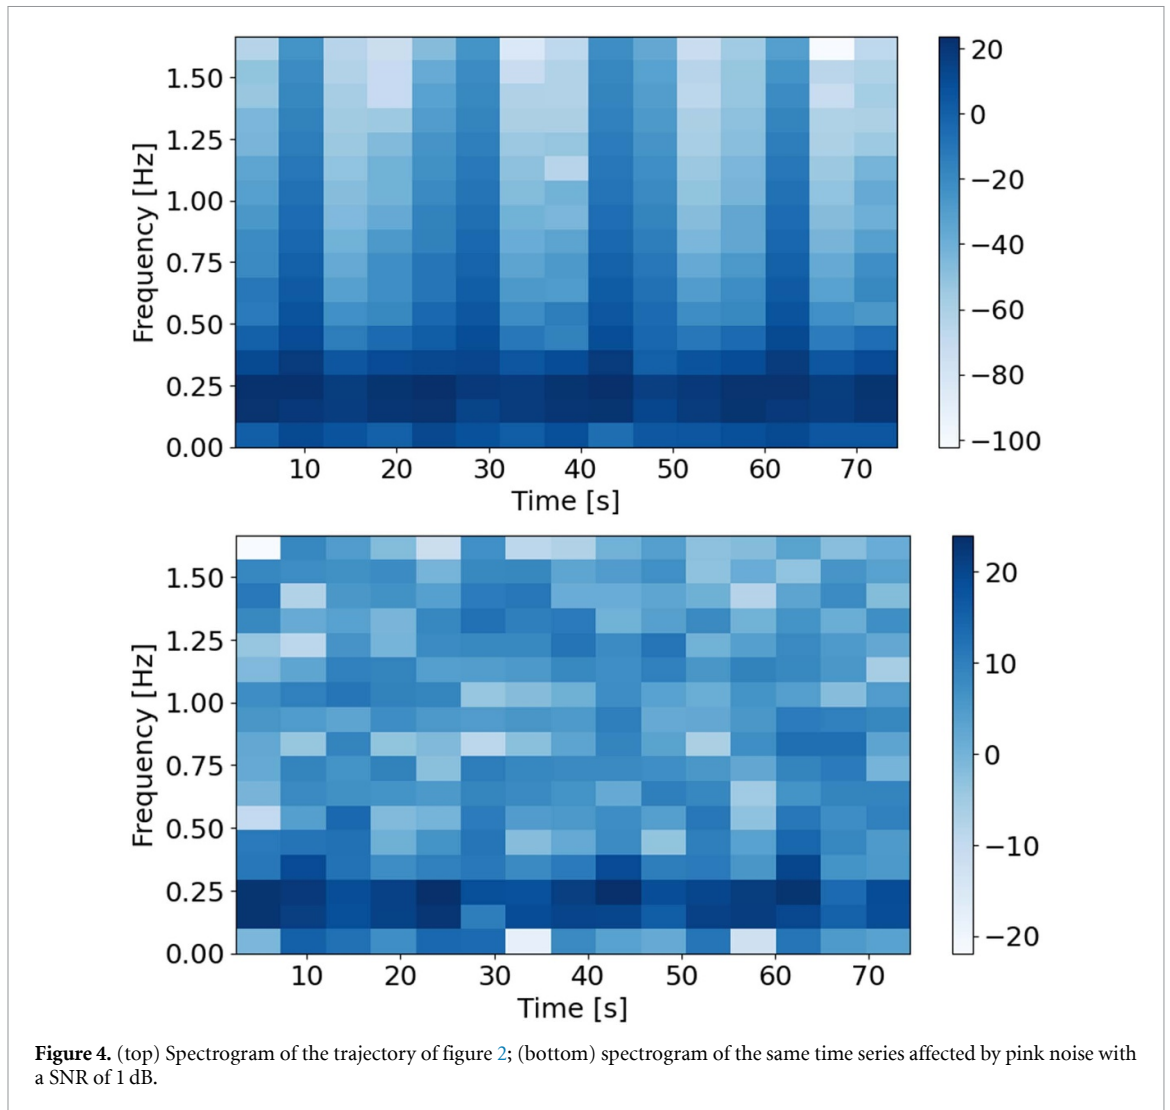
### 3.5. Multimodal DNNs

One of the most significant breakthroughs of deep learning that enabled DNN to reach unprecedented performance in solving complex tasks is the capability of DNNs to conduct their analysis directly on raw, un-processed, data. This peculiarity of DNNs is caused by their ability to learn through the learning process hierarchical representations of their training data, hence conducting an automatic feature extraction process [29].
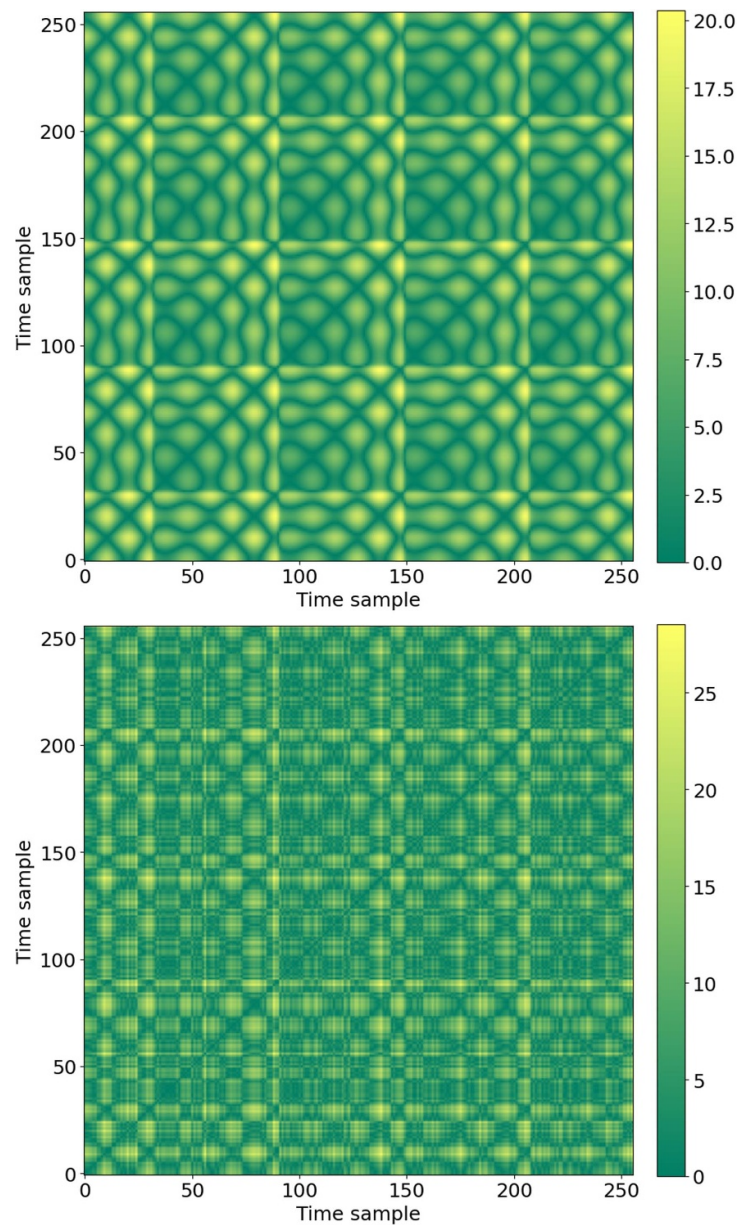
**Figure 3.** Unthresholded recurrence plot of the time series of figure 2(top) and of the same time series affected by pink noise with a SNR of 1 dB (bottom).

Despite the well-established results of DNN in terms of automated feature extraction, a significant effort was spent on designing specialized DNN architectures able to capture more easily certain features and patterns in different data structures (e.g. convolutional DNN to recognize geometrically-close patterns in images). In this direction, multimodal DNN [30] are a class of DNN specialized to deal with heterogeneous input data, that is, data of various structures/*modalities* (e.g. the combination of images, signal readings and text). In multimodal DNNs, the first layers of the DNN are replaced by multiple separate *channels* that analyze individually and in parallel the various inputs. This separate processing allows the employment of different specialized architectures on the various channels, hence carrying out a more effective feature extraction on each modality.

The extracted features are then combined in the deeper layers of the DNN, allowing for a more comprehensive, yet less complex, analysis of the heterogeneous data in its entirety.

**Figure 4.** (top) Spectrogram of the trajectory of figure 2; (bottom) spectrogram of the same time series affected by pink noise with a SNR of 1 dB.

## 4. Proposed multimodal DNN

In this section we detail the multimodal DNN architecture and its training process.

### 4.1. DNN architecture

The proposed multimodal architecture is depicted in figure 5 and consists of three separate *convolutional channels* that converge into a common set of *fully connected* (FC) layers, which ends with a 15-neuron layer which implement a softmax function to perform the classification among the 15 different dynamics assumed in section 3.1.

FC layers [31] are the simplest element of DNN architectures and are formed by a set of neurons, each connected to every neuron of the previous layer. Having all possible connections, FC layers have a very high capacity (i.e. the degree of complexity of the function they can approximate). Hence, in complex architectures, FC are typically used as final layers to produce the DNN output based on the results obtained by some more specialized layers.

Convolutional layers [31] are the constituting elements of convolutional neural networks and exploit the local connectivity of their neurons, which share the same weights, to implement a sliding-window/convolutional analysis of their input. Due to the geometric nature of the convolution process, convolutional layers have found vast applications in vision-related tasks and signal processing, serving the purpose of extracting information (or *features*) from an image/signal-like input while preserving the geometric/spatial nature of the data. Convolutional layers are hence typically used as the first layers of a DNN as *feature extractors* for the deeper layers, which are commonly FC.

In our architecture, we employ three different convolutional channels to analyze separately the three different inputs (i.e. a time series, its recurrence plot and its spectrogram). For the recurrence plot and spectrogram channels, since their input consists of images, we used 2D-convolutional layers, whereas
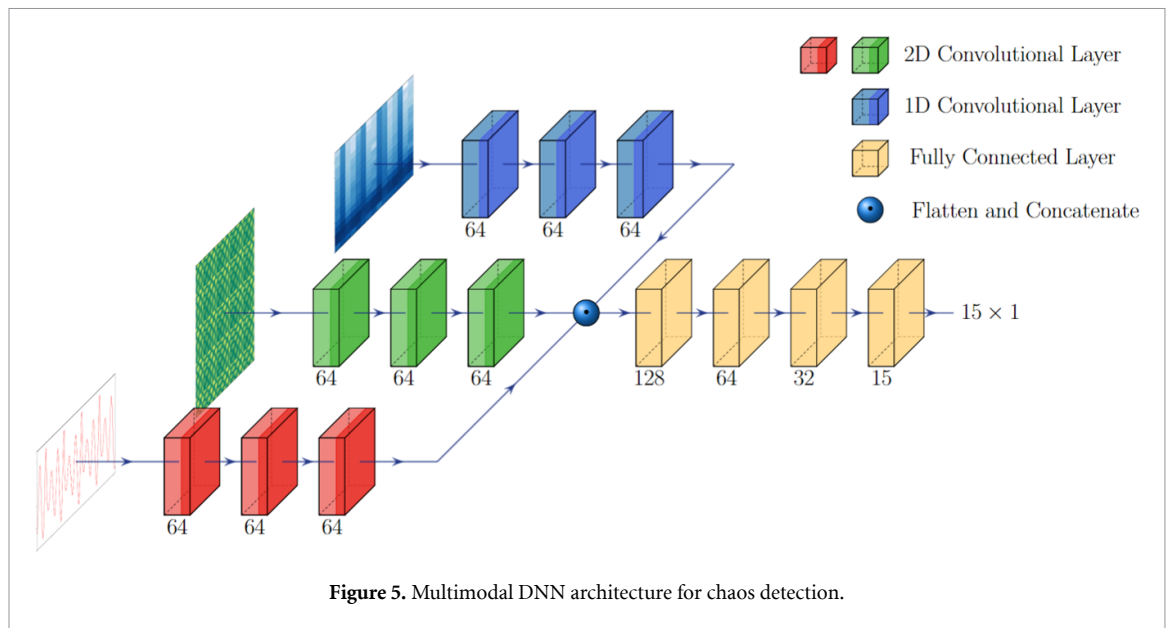
**Figure 5.** Multimodal DNN architecture for chaos detection.

1D-convolutions were used for the time series channel. The output of the three convolutional channels is then flattened and concatenated before being fed to the final, common, FC portion of the DNN. For the sake of completeness, we mention that pooling layers [31], which are the third element commonly found in standard convolutional DNNs such as LeNets [32] and VGG networks [33], were not included in our architecture. In fact, the main advantage of using max/average pooling layers is to reduce the dimension of the feature maps internally processed by the convolutional neural network by compressing information and, in turn, lower the number of trainable parameters used. Given the reasonably limited number of parameters involved in the training phase of our DNN, which is kept low also thanks to the multi-step training procedure detailed in section 4.2, we did not observe any significant benefit from their inclusion but, in principle, they may be seamlessly included in deeper networks for chaos detection derived from ours.

We remark that, since both the recurrence plots and spectrograms are derived from the measurement time series, the transformations needed to obtain them can be seen as initial feature extractors for the time series, that are further refined by the CNN channels.

For our tests, we set all neuron activation functions to be rectified linear units (ReLU), save for the final layer that used the Softmax activation function as customary in classification tasks.

### 4.2. Training process

To better capture the characteristics of each data representation, the three channels were trained separately and then combined in a multi-step training procedure. This approach reduces the resources needed for the training, as only a portion of the DNN is trained at a given time, but its main advantage is that it ensures that each convolutional channel produces informative feature maps to be fed to the final FC layers. In fact, since such feature maps were originally used individually to solve the same classification task, their usage allows for a simpler training and reduces the tendency of multimodal DNNs to rely mostly on a single modality [34].

To train the convolutional layers of the three channels represented in figure 5, we first trained three separate DNNs, each constituted by such convolutional layers followed by three additional FC layers with 64, 32 and 15 neurons. This individualized training requires the three DNNs to recognize the chaotic dynamics relying only on the single data modality provided (e.g. only on the recurrence plot), meaning that the three convolutional layer stacks become feature extractors specialized for their respective data.

The FC layers are then discarded, and the trained convolutional layers are placed in the multimodal architecture depicted in figure 5. The multimodal network is then further trained, this time using all data modalities at the same time, keeping *frozen* the previously obtained weights of the convolutional layers and allowing the update of only the layers following the concatenation, as it is done in a transfer learning procedure.

The presented training process is one of the possible approaches to train a multimodal model designed following the *late fusion* paradigm [35–37], as is the case for our DNN. Compared to the others (e.g. ensembling multiple single-modal DNNs) it provides a solution that correlates/fuses the various modalities through an ad-hoc learning process (that is, the training of the final FC layers), while avoiding the more complex task of training directly the multimodal DNN in its entirety, which may be affected by a so-called

*greedy learner* behavior (that is a tendency to rely on a single modality) and may require significantly more hyperparameters tuning. The proposed multi-step training procedure hence allows for better information fusion regarding the analysis of the three data modalities, leading to an overall better performance when compared to single-channel DNNs, as the one used in [15], or multi-model architectures trained directly on multiple input types (as it will be shown in section 5). Furthermore, we stress that the proposed procedure requires significantly less memory and computing power when compared to the direct training of the entire multimodal DNN, making the system more scalable and allowing to easily expand the DNN with additional channels if needed.

# 5. Simulations

### 5.1. Dataset description

As mentioned, the dataset used for our testing has been derived from the one presented in [3] and made available by the authors in [26]. The starting dataset included time series sampled from a total of 15 different dynamical systems, including 5 chaotic systems and 10 systems generating periodic, quasi-periodic and non-periodic signals. For our tests, we considered 1000 trajectories for each of the 15 systems, generating additional ones when needed using the original code from the authors made available in [26], and reserved 25% of data for validation.

As mentioned, we extracted from the dataset only the trajectories reporting the first component of the various systems, from which we extracted time series of 256 samples. Such trajectories were then corrupted by additive random noise so that the resulting signal-to-noise-ratio (SNR) of each trajectory was respectively equal to 5 dB and 1 dB for our first and second tests. Regarding the noise nature, each trajectory was subject with a uniform probability to either white, pink or brown noise, generated using the colored noise python library [38]. Spectrograms were constructed using the specgram() function from the matplotlib library using 32 points for the fast Fourier transform blocks with 16 points of overlap.

All training was conducted on a machine equipped with an RTX 3090 GPU. All of the three convolutional networks were trained for 15 epochs, which required about 30 s each, with an additional 15 epochs reserved for the fine-tuning of the last layers of the multimodal DNN, in line with the procedure of section 4.2.

To identify our classes, we will follow the naming scheme of [3] that uses the following suffixes: CHA = chaotic; QPS = quasi-periodic; OSC/DOSC/IOSC = undamped/damped/rising oscillators; DS = damped systems. Classes were then numbered according to the naming in [3] as: 0 = OSC_1; 1 = DS_1; 2 = OSC_2; 3 = CHA_4 (an autonomous dissipative flow showing a Halvorsen attractor); 4 = QPS_2; 5 = CHA_5 (an autonomous dissipative flow showing a Rucklidge attractor); 6 = QPS_1; 7 = CHA_3 (an autonomous dissipative flow showing a Rössler attractor); 8 = DOSC_1; 9 = DOSC_2; 10 = CHA_2 (the Lorenz system); 11 = QPS_3; 12 = IOSC; 13 = CHA_1 (a driven dissipative flow showing a Ueda oscillator); 14 = DS_2.

### 5.2. Results

Figure 6 reports the validation confusion matrix, normalized by rows between 0 and 1, of our first test in which we set the SNR to 5 dB. The resulting total accuracy (number of successful predictions made by the model over all the classes divided by the total number of predictions) is about 94%, with most miss-classifications occurring for class 1 that is mistaken for class 14. These mistakes can be explained by considering that classes 1 and 14 both represent damped linear systems that differ mostly by the magnitude of their damping. Figure 7 confirms the good quality of the classification by reporting the F1-score archived by our DNN: on the right side of the figure, every class is given its own class F1-score computed as the harmonic mean of the class precision and recall; on the left portion of the figure, we report the macro-F1 score (that is the arithmetic average of the class F1-scores), the micro-F1 score (that is the harmonic mean of the overall precision and recall of the model) and the weighted-F1 score (that is the arithmetic average of the class scores weighted by their sample number).

For the sake of comparison and to validate the effectiveness of our multi-step training procedure described in section 4.2, we mention that the best-performing single-modal DNN was the one trained on RPs which had an overall accuracy of about 85%, whereas training directly the entire multimodal architecture of figure 5 led to an overall accuracy of about 90%.

For our second test, we significantly increase the noise power, reaching an SNR of 1 db. Figure 8 highlights how the DNN accuracy drops for all classes, reaching an average accuracy of about 76% (65% for single channels), with the most affected classes being once again 1 and 14. Class 9 is also commonly mistaken for class 8 (both are oscillating dynamics with decreasing amplitude), whereas class 4 (a quasi-periodic nonlinear system) is mistaken mostly for the oscillating dynamics of class 8 or for class 6 (that is another quasi-periodic nonlinear system). Overall, we can conclude that as expected the higher level of noise makes the distinction between similar sets of dynamics more complex; however, chaos detection still reaches

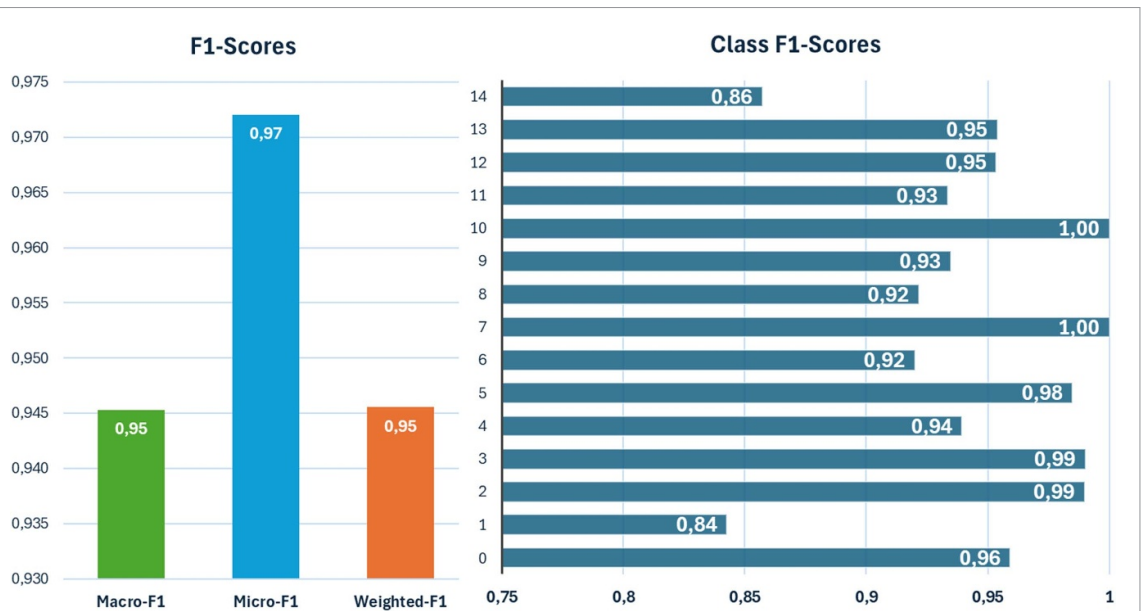**Figure 6.** Confusion matrix for the first test (SNR = 5 dB).



**Figure 7.** F1-Scores for the first test (SNR = 5 dB).

satisfactory results with the chaos-related classes (3, 5, 7, 10, 13) maintaining an adequate level of performance, as also supported by the F1-scores of figure 9.

As a final test, we task our DNN with classifying trajectories sampled from new, unseen, dynamics that were not included in the training set. To do so, we generated a new dataset of 1000 trajectories for each of the five dynamics of our chaotic systems. Each of such trajectories was generated by randomly perturbing each of the parameters of the original dynamics of a factor equal to $\pm 5\%$ of its original values. Figure 10 displays the confusion matrix of such a classification with a snr equal to 1 dB: from the diagonal entries of the matrix it is

**Figure 8.** Confusion matrix for the second test (SNR = 1 dB).



**Figure 9.** F1-Scores for the second test (SNR = 1 dB).

**Figure 10.** Confusion matrix for the third test (only chaotic systems with different parameters).

clear that the perturbed dynamics are harder to recognize, however overall the DNN maintains an accuracy of 82.2%, with a micro-F1 score of 0.90 and macro-F1 score of 0.86, suggesting that the DNN is capable of identifying to a reasonable extend which known dynamics is the most similar to the one observed.

## 6. Conclusions and future works

This work has presented the design of multimodal DNN for the problem of chas detection in measurements affected by noise. The proposed DNN was designed to solve a multi-class classification problem, with the goal of not only detecting chaotic behavior but also identifying the specific attractor that characterizes the underlying system dynamics. The ad-hoc architecture of the multimodal DNN we designed involves three parallel input channels, each respectively specialized for the analysis of recurrence plots, spectrograms and time series. The overall DNN was trained following a custom multi-step procedure that consists in first training each channel separately and then training the common portion of the DNN to correctly fuse the information extracted from the various input modalities. Our system was able to discern among 15 different dynamics from the measurement of a single state component, demonstrating high accuracy in discerning chaotic dynamics from non-chaotic ones even in the presence of significant levels of noise.

Future works are related to increasing the number of dynamics considered by extending the dataset with some other relevant chaotic systems and in extending the capabilities of the proposed DNN. Regarding the multimodal input, we aim at adding more complex and informative inputs/channels into the DNN, whereas regarding the output we are designing a DNN that employs multiple output channels to provide different analysis: a first relevant output to include would be a neural operator capable of estimating the Lyapunov exponent [8], while a more challenging and innovative study would integrate an output channel consisting of a Kolmogorov–Arnold network [39] to provide a symbolic guess of the observed dynamics.

## Data availability statement

No new data were created or analysed in this study.

## Acknowledgments

## ORCID iD

Alessandro Giuseppi ● https://orcid.org/0000-0001-5503-8506

## References

[1] Gottwald G A and Melbourne I 2009 On the implementation of the 0–1 test for chaos *SIAM J. Appl. Dyn. Syst.* **8** 129–45
[2] Toker D, Sommer F T and D'Esposito M 2020 A simple method for detecting chaos in nature *Commun. Biol.* **3** 11
[3] Szczesna A, Augustyn D, Harezlak K, Josinski H, Switonski A and Kasprowski P 2023 Datasets for learning of unknown characteristics of dynamical systems *Sci. Data* **10** 79
[4] Skiadas C H and Skiadas C 2017 *Handbook of Applications of Chaos Theory* (CRC Press)
[5] Basar E 1990 *Chaos in Brain Function* (Springer)
[6] Tito Arecchi F and Harrison R G 2012 *Instabilities and Chaos in Quantum Optics* vol 34 (Springer)
[7] Kumar A 2012 Chaos theory: impact on and applications in medicine *J. Health Allied Sci. NU* **02** 93–99
[8] Abarbanel H D, Brown R and Kennel M B 1991 Lyapunov exponents in chaotic systems: their importance and their evaluation using observed data *Int. J. Mod. Phys.* B **05** 1347–75
[9] Wolf A, Swift J B, Swinney H L and Vastano J A 1985 Determining lyapunov exponents from a time series *Physica D* **16** 285–317
[10] Gottwald G A and Melbourne I 2016 *The 0–1 Test for Chaos: A Review* (Springer) pp 221–47
[11] Barrio R, Lozano A, Mayora-Cebollero A, Mayora-Cebollero C, Miguel A, Ortega A, Serrano S and Vigara R 2023 Deep learning for chaos detection *Chaos* **33** 073146
[12] Boullé N, Dallas V, Nakatsukasa Y and Samaddar D 2020 Classification of chaotic time series with deep learning *Physica D* **403** 132261
[13] Seok Lee W and Flach S 2020 Deep learning of chaos classification *Mach. Learn.: Sci. Technol.* **1** 045019
[14] Zanin M 2022 Can deep learning distinguish chaos from noise? numerical experiments and general considerations *Commun. Nonlinear Sci. Numer. Simul.* **114** 106708
[15] Mukhopadhyay S and Banerjee S 2020 Learning dynamical systems in noise using convolutional neural networks *Chaos* **30** 103125
[16] Hassona S, Marszalek W and Sadecki J 2021 Time series classification and creation of 2d bifurcation diagrams in nonlinear dynamical systems using supervised machine learning methods *Appl. Soft Comput.* **113** 107874
[17] Rappeport H, Levin Reisman I, Tishby N and Balaban N Q 2022 Detecting chaos in lineage-trees: a deep learning approach *Phys. Rev. Res.* **4** 013223
[18] Wenkack Liedji D, Hervé Talla Mbé J and Kenné G 2022 Chaos recognition using a single nonlinear node delay-based reservoir computer *Eur. Phys. J.* B **95** 18
[19] Fan H, Jiang J, Zhang C, Wang X and Lai Y-C 2020 Long-term prediction of chaotic systems with machine learning *Phys. Rev. Res.* **2** 012080(R)
[20] Lahmiri S and Bekiros S 2019 Cryptocurrency forecasting with deep learning chaotic neural networks *Chaos Solitons Fractals* **118** 35–40
[21] Wang R, Kalnay E and Balachandran B 2019 Neural machine-based forecasting of chaotic dynamics *Nonlinear Dyn.* **98** 2903–17
[22] Thavarajah R, Zhai X, Zheren M and Castineira D 2021 Fast modeling and understanding fluid dynamics systems with encoder–decoder networks *Mach. Learn.: Sci. Technol.* **2** 025022
[23] de Jesús Serrano-Pérez J, Fernández-Anaya G, Carrillo-Moreno S and Wen Y 2021 New results for prediction of chaotic systems using deep recurrent neural networks *Neural Process. Lett.* **53** 1579–96
[24] Sangiorgio M, Dercole F and Guariso G 2021 Forecasting of noisy chaotic systems with deep neural networks *Chaos Solitons Fractals* **153** 111570
[25] Zhou S, Zhao Z and Wang X 2022 Novel chaotic colour image cryptosystem with deep learning *Chaos Solitons Fractals* **161** 112380
[26] Augustyn D 2023 Dataset for learning of unknown characteristics of dynamical systems - code (available at: https://gitlab.com/draugustyn/signal-data)
[27] Marwan N, Carmenromano M, Thiel M and Kurths J 2007 Recurrence plots for the analysis of complex systems *Phys. Rep.* **438** 237–329
[28] Iwanski J S and Bradley E 1998 Recurrence plots of experimental data: to embed or not to embed? *Chaos* **8** 861–71
[29] Wiatowski T and Bolcskei H 2018 A mathematical theory of deep convolutional neural networks for feature extraction *IEEE Trans. Inf. Theory* **64** 1845–66
[30] Ramachandram D and Graham W T 2017 Deep multimodal learning: a survey on recent advances and trends *IEEE Signal Process. Mag.* **34** 96–108
[31] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (MIT Press)
[32] Lecun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition *Proc. IEEE* 86 2278–324
[33] Simonyan K and Zisserman A 2014 Very deep convolutional networks for large-scale image recognition (arXiv: 1409.1556)
[34] Nan W, Jastrzebski S, Cho K, and Geras K J 2022 Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks. *Proc. 39th International Conf. on Machine Learning* (*Proc. Machine Learning Research*), vol 162, ed K Chaudhuri, S Jegelka, L Song, C Szepesvari, G Niu, and S Sabato, pp 24043–55 (PMLR)
[35] Liu K, Yanen Li, Ning X and Natarajan P 2018 Learn to combine modalities in multimodal deep learning (arXiv: 1805.11730)

[36] Gadzicki K, Khamsehashari R and Zetzsche C 2020 Early vs late fusion in multimodal convolutional neural networks *2020 IEEE 23rd International Conf. on Information Fusion (FUSION)* (IEEE)

[37] Eitel A, Tobias Springenberg J, Spinello L, Riedmiller M and Burgard W 2015 Multimodal deep learning for robust rgb-d object recognition *2015 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)* (IEEE)

[38] Burkardt J. Colored noise 1/f power law noise generation (available at: https://people.math.sc.edu/Burkardt/cpp_src/colored_noise/colored_noise.html)

[39] Liu Z, Wang Y, Vaidya S, Ruehle F, Halverson J, Soljačić M, Hou T Y, and Tegmark M 2024 Kan: Kolmogorov-Arnold networks (arXiv: 2404.19756)