



Full length article



A novel Out-of-Distribution detection approach for Spiking Neural Networks: Design, fusion, performance evaluation and explainability

Aitor Martinez-Seras^{a,b,*}, Javier Del Ser^{a,c}, Jesus L. Lobo^a, Pablo Garcia-Bringas^b, Nikola Kasabov^{d,e,f}

^a TECNALIA, Basque Research and Technology Alliance (BRTA), 48160 Derio, Spain

^b University of Deusto, 48007 Bilbao, Spain

^c University of the Basque Country (UPV/EHU), 48013 Bilbao, Spain

^d Auckland University of Technology, Auckland, New Zealand

^e Intelligent Systems Research Center, Ulster University, UK

^f IICT, Bulgarian Academy of Sciences, Bulgaria

ARTICLE INFO

Keywords:

Spiking Neural Networks
Out-of-Distribution detection
Explainable artificial intelligence
Model fusion
Relevance attribution

ABSTRACT

Research around Spiking Neural Networks has ignited during the last years due to their advantages when compared to traditional neural networks, including their efficient processing and inherent ability to model complex temporal dynamics. Despite these differences, Spiking Neural Networks face similar issues than other neural computation counterparts when deployed in real-world settings. This work addresses one of the practical circumstances that can hinder the trustworthiness of this family of models: the possibility of querying a trained model with samples far from the distribution of its training data (also referred to as Out-of-Distribution or OoD data). Specifically, this work presents a novel OoD detector that can identify whether test examples input to a Spiking Neural Network belong to the distribution of the data over which it was trained. For this purpose, we characterize the internal activations of the hidden layers of the network in the form of spike count patterns, which lay a basis for determining when the activations induced by a test instance is atypical. Furthermore, a local explanation method is devised to produce attribution maps revealing which parts of the input instance push most towards the detection of an example as an OoD sample. Experimental results are performed over several classic and event-based image classification datasets to compare the performance of the proposed detector to that of other OoD detection schemes from the literature. Our experiments also assess whether the fusion of our proposed approach with other baseline OoD detection schemes can complement and boost the overall OoD detection capability. As the obtained results clearly show, the proposed detector performs competitively against such alternative schemes, and when fused together, can significantly improve the detection scores of their constituent individual detectors. Furthermore, the explainability technique associated to our proposal is proven to produce relevance attribution maps that conform to expectations for synthetically created OoD instances.

1. Introduction

In the last decades, Machine Learning (ML) has been widely proven to outperform human capabilities in tasks that were previously thought to be intractable for machines, such as image/video classification, natural language processing or generative modeling. The unprecedented levels of performance achieved over these tasks have spurred an ever-growing proliferation of practical applications exploiting the capabilities of ML models. The sensitive nature of practical decisions that can be made as per the output elicited by ML models has lately motivated a flurry of research around the need for ensuring the trustworthiness of

the audience consuming their inputs. This issue is particularly concerning when it comes to models whose internal structure and/or learning algorithm cannot be easily understood by non-expert users. This is the case of Deep Learning, arguably the field where the most outstanding breakthroughs have been reported in manifold domains [1,2]. Concerns with the trustworthiness of Deep Learning models are even more acute as a result of their inherent black-box nature. Consequently, a new field of study has developed over the last decade, coined as eXplainable Artificial Intelligence (XAI), which focuses on making the decisions issued by ML models understandable for humans [3–5].

* Correspondence to: Parque Tecnológico de Bizkaia, 700, 48160 Derio, Bizkaia, Spain.
E-mail address: aitor.martinez@tecnalia.com (A. Martinez-Seras).

<https://doi.org/10.1016/j.infus.2023.101943>

Received 17 March 2023; Received in revised form 5 July 2023; Accepted 21 July 2023

Available online 26 July 2023

1566-2535/© 2024 Fundación Tecnalia Research & Innovation. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Besides the lack of interpretability of ML models, other factors may hinder the trustworthiness of ML models when deployed in real-world contexts. In such scenarios, models can face more varying and complex circumstances than the ones encountered in the controlled environments where their parameters are usually learned. Among them, one of the essential assumptions made in ML is that both training and test data are drawn from the same distribution. However, this assumption may not hold in real-world settings, in which this distribution may change, and test instances might be sampled from a unknown distribution. In general ML models are not designed from scratch to discriminate between *known* and *unknown* query samples as per the knowledge captured during their training processes. In this case, it becomes necessary to endow the model with the ability to determine whether the queried sample belongs to the data distribution it has learned. This task is exactly what Out-of-Distribution (OoD) detection addresses [6]: unlike outlier detection, OoD methods depart from a model already trained over a given dataset, namely, the *In-Distribution* (ID) dataset. The goal pursued in OoD detection is to declare whether a new test instance has either been sampled from the modeled ID dataset or, instead, it belongs to another distribution (correspondingly, the *Out-Distribution* dataset, OD). Given its straightforward utility in uncontrolled environments, OoD detection has been acknowledged to be of pivotal importance in many applications, including safe autonomous driving [7] and medical diagnosis [8,9].

From a modeling perspective, the literature has tackled the OoD detection problem by focusing mainly on neural networks for image classification, using convolutional architectures for the purpose. In general, OoD instances can be identified by adopting (i) model-agnostic detection methods; or (ii) model-specific strategies. The former usually operate by characterizing information elicited by the model during its training process that is not particular of the model under target (e.g. the produced output class distribution of training samples). Model-specific methods, however, inspect the internal parameters of the model where its captured knowledge is persisted (e.g. forward-propagated activations of neurons). Above all, most detectors reported so far by the community remain as opaque as the models they aim to supervise, missing to attribute which parts of the input test instance are most influential for the detection of the whole sample as an OoD instance. In other words, OoD detection must be also explained for the audience of the model; otherwise, it could jeopardize the completion of a fully trustworthy ML pipeline. Another gap in the current literature is the lack of model-specific proposals suitable to detect OoD instances in Spiking Neural Networks (SNNs). These models, often referred to as *the third generation of artificial neural networks*, are the evolution of present-day neural networks in virtue of their energy efficiency when implemented in specialized neuro-morphic hardware, and their ability to model complex temporal dynamics thanks to their event-based nature [10,11]. The latter, makes SNNs the best suited models to process event-based data like data coming from the Dynamic Vision Sensor (DVS), i.e., a bio-inspired vision sensor or camera that outputs pixel-level brightness changes instead of standard intensity frames, offering advantages like a very high dynamic range, no motion blur, and a latency in the order of microseconds [12]. Despite the intense research activity around this family of models noted in recent times, no prior work exists dealing with the detection of OoD instances by leveraging specifics of the spike-based working procedure of SNNs.

This work covers the gap identified above in OoD research by developing a novel model-specific detection technique specifically designed for *fully-spiking* SNNs (i.e. neural architectures that exclusively rely on spikes for their internal activations). Our proposed approach utilizes information produced by the SNN learning algorithm during the training process to build a representation of its typicality when processing ID data. A second contribution of our work is a local attribution method that highlights the parts of the input instance that push the developed detector towards classifying a sample as OoD. The attribution information produced by this method can improve

the acceptability of the audience of the model when being notified about an *unknown* input, connecting clearly to the trustworthiness of ML pipelines sought in risk-sensitive domains as the ones exemplified previously. An extensive experimental setup over several image classification tasks is designed (i) to compare the devised OoD detection method to existing alternatives that can be adapted to SNNs; (ii) to verify whether a fusion of baseline OoD detection techniques with our proposed detector can be synergistic and attain even better detection scores than those achieved by the detectors on their own; and (iii) to assess whether the generated attributions succeed at identifying artificially induced image artifacts on the considered datasets. Results are promising, exposing a very competitive detection performance and attributions that conform to expectations.

The rest of the paper is structured as follows: first, Section 2 contextualizes the OoD detection task, formally stating the problem it aims to solve, and introducing a taxonomy that allows framing the contribution of the work. This section also briefly mentions recent advances made in SNNs and XAI. Next, Section 3 presents the proposed detector by describing its overall workflow, followed by a detailed description of the detection algorithm and the procedure to extract local explanations in the form of attribution maps. The experimental setup is described in Section 4, whereas results are discussed quantitatively and qualitatively in Section 5. Finally, conclusions are drawn in Section 6, together with several research paths rooted on our investigations.

2. Related work

This study lies in the intersection between several fields of research. Consequently, this section briefly pauses at the state of the art of each of such fields: the OoD detection problem (Section 2.1), Spiking Neural Networks (Section 2.2) and Explainable AI in SNNs (Section 2.3). Finally, we end by stating the contribution of this work beyond the current status of knowledge in these surveyed areas (Section 2.4).

2.1. Out-of-distribution detection

Before formulating the problem in a mathematical fashion, it is important to clearly define what the out-of-distribution detection encompasses and to establish the similarities and differences with other related tasks, such as outlier, anomaly and novelty detection. This is precisely the central matter tackled in the comprehensive survey recently contributed in [6], which also points out several misconceptions existing in the literature. This work distinguishes between two different types of shift in data: the *semantic shift* and the *covariate shift*. The latter stands for the situation where features of the input data vary without any change in the class they belong to. Stated differently, the statistical dependence between variables change, but not their relationship with the output variable. In contrast, semantic shift refers to the case when the new data belongs to a class different than the ones learned by the model during its training stage. This being said, it is important to note that OoD detection addresses only semantic shift (see Fig. 1).

To clarify the application domain in which OoD detection can be formulated, we start by differentiating it from the above-mentioned tasks. The most distinct task is outlier detection, as it does not follow the classic train-test schema. Instead, a set of observations is provided, and the task is to detect significantly different samples from this contaminated dataset of observations. Both semantic and covariate shifts can occur in this scenario. Conversely to outlier detection, OoD detection is assumed to occur under a train-test schema. As concluded in [6], OoD detection can be regarded as a super-category that includes semantic anomaly detection, multi-class novelty detection and open set recognition [13]. Summarizing, in OoD detection a model is learned from a training set that is sampled from a given ID dataset, whereas test samples are drawn evenly from either the in-distribution or from a semantically different distribution, namely, the aforementioned OD dataset. The objective is to correctly classify test samples as belonging




Fig. 1. Difference between semantic and covariate shift. The semantic one occurs when a new class (star) emerges. In the covariate shift, the features of the input change (the color of the figures, from red to blue) while maintaining the semantics (the label), what makes the model incorrectly classify some instances.

to the ID or OD distributions, while maintaining the performance of the model in the original classification task. In almost all cases, detection is ultimately achieved by crafting a scoring function that assigns different scores for the in- and out-distribution samples.

Mathematically, the OoD detection task aims to distinguish whether a query instance \mathbf{x}_{query} belongs to the data distribution of the data used to train the model, or alternatively can be thought to belong to another different distribution. Without loss of generality, let the *in-distribution* dataset be denoted as D_I , which is drawn from a probability distribution $P_X(\mathbf{x})$ defined over a space $\mathcal{X} \in \mathbb{R}^K$. This dataset is used to learn a classifier $M_\theta(\mathbf{x})$ by adjusting the model's parameters θ so as to maximize its generalization capability over a set of labels $\{1, \dots, \ell\}$. Once the model has been learned, the classifier can infer the class $\hat{y} = M_\theta(\mathbf{x}) \in \{1, \dots, \ell\}$ of any new sample $\mathbf{x} \in \mathbb{R}^K$. We further assume another distribution $Q_X(\mathbf{x})$ different than $P_X(\mathbf{x})$, yet over the same space \mathcal{X} (*out-distribution*). New query samples can be sampled from a mixture distribution $M_{X|Z}(\mathbf{x}|z)$ such that $M_{X|Z}(\mathbf{x}|z)$ is equal to $P_X(\mathbf{x})$ if $z = 1$, and $Q_X(\mathbf{x})$ if $z = 0$. If Z follows a uniform binary probability distribution, an OoD detector $G_\phi(\mathbf{x})$ with parameters ϕ aims to distinguish which test instances follow the in-distribution from the ones which not, without observing variable Z , and by only resorting to the trained model $M_\theta(\mathbf{x})$.

Recalling to [6] and following its proposed taxonomy, OoD detection methods can be divided in three types: classification-based, density-based or distance-based. Our technique falls within the latter category, as it comprises a detector that relies on the computation of a measure of distance between the query samples and the centroids or prototypes of the classes present in the in-distribution. Nevertheless, we herein focus on the literature related to some of the classification-based methods, specifically those which can be labeled as *post-hoc* methods. These detectors can be easily adapted to Spiking Neural Networks without making changes to their core, by solely resorting to information produced at their output (e.g., the logits).

The first paper coining the *Out of Distribution detection* term was actually a classification-based method, presented in [14] and often referred to as the *baseline* method. This approach utilizes the so-called Maximum Softmax Probability (MSP) to detect OoD samples, based on the observation that samples belonging to the ID dataset tend to have a higher MSP value than OD instances. By simply defining a threshold on this score, their method achieved acceptable results in a wide variety of datasets. The authors in [15] went one step further by applying a temperature scaling strategy to the softmax computation, which pushes softmax scores of in- and out-distribution samples further apart from each other. In addition, they included input preprocessing by adding small gradient and softmax dependent perturbations to the data, slightly improving the performance. More recently, the work in [16] assumed that the feature space of the penultimate layer of a classifier follows a Gaussian distribution, allowing for the estimation of mean and variance statistics from the features of every class. A posterior fit of a multivariate Gaussian distribution and the usage of Mahalanobis distance to gauge the closest class-conditional distribution is used for OoD detection. Following the strategy posed in [15], the performance is improved by adding small noise perturbations; this time depending on the Mahalanobis distance rather than softmax. Further rationale for

the effectiveness of Mahalanobis distance for OoD detection is given in [17].

Contrarily to previous approaches, in [18] a technique called *Outlier Exposure* is presented for improving the performance of existing detectors. Authors propose to leverage the enormous quantity of data available nowadays by modifying the training process of neural networks with the addition of an additional term to the original loss function. This term depends on the original task (classification, density estimation, etc.) and on the detector in use. This term helps the model to learn heuristics that will improve the detector's performance. Later, an energy-based detector was derived in [19] by adapting the concept of *Helmholtz free-energy* to deep neural networks, expressing it in terms of the denominator of the softmax activation. Therefore, each sample is assigned a score based on its logits (its *energy*), which is shown to be higher for OoD samples than for ID instances. This was proven to yield even better detection performance than previous softmax-based scores. Thereupon, the study presented in [20] improved the computational efficiency at inference time of the energy-based detector by training multiple intermediate classifiers operating at different depths of the trained neural network. The intermediate activations of the neural networks are extracted at certain depths of its internal layered structure depending on its complexity, leveraging the intuition that less complex OoD samples can be detected by using low-level activation statistics of the network. Extending the success of the energy-based detectors, authors in [21] came up with a new scoring function (*JointEnergy*), which exploits the joint uncertainty across labels by aggregating the label-wise free energy for multiple labels. In contrast to presented methods, *GradNorm* [22] explores the model's parameters and the gradient space to compute an OoD score. Specifically, GradNorm backpropagates the Kullback-Leibler (KL) divergence between the softmax output and the uniform distribution, computing the OoD score by extracting the gradients of the last fully connected layer and calculating its L1 norm. Recently, [23] proposed a simple post-hoc method forged as *Rectified Activations* (ReAct), which is applied to the penultimate layer of a network to truncate the activations at inference time. Its effectiveness is supported by the finding that applying the BatchNorm statistics of ID training data to the OoD data at inference time leads to abnormally high activations in all layers, and therefore to the model's output, deteriorating the detection performance. Basically, ReAct limits the activation's value, yielding better detection scores than the previously presented techniques hinging on the logits of the network. One of the most recently proposed OoD methods is *RankFeat* [24], which relies on subtracting the rank-1 matrix from the high-level feature, and using logits derived from this perturbation to compute the OoD score. They demonstrate that in-distribution samples are less affected by this perturbation, enabling the detection of OoD samples in the logit space.

2.2. Spiking neural networks

This family of neural networks has garnered significant attention since their inception, not only because they can represent more faithfully the inner stimuli of the human brain, but also due to the technical advantages that their particular information processing strategy entails in terms of spatio-temporal pattern learning and energy efficiency. In

the brain, information flowing between neurons is conveyed through synapses by trains of short electric pulses called *spikes*, so that neurons receive and accumulate them as potential in their membrane. When a potential threshold is met, the neuron emits a pulse (spike) of its own to the subsequent neuron. The neuron emitting a spike is the pre-synaptic neuron, whereas the receiver neuron is referred to as post-synaptic neuron. Mimicking this dynamic behavior offers an inherent way to deal with temporal data, and the sparsity and binary nature of spikes makes it possible to reduce dramatically the energy consumption of these models when deployed on specialized hardware, as chips like IBM's TrueNorth [10], The SpiNNaker Chip [25] and Intel's Loihi [26].

A major drawback of SNNs is the fact that spiking neurons are not differentiable due to the discontinuity of the spike. Therefore, gradient backpropagation-based learning schemes are not directly applicable. To overcome this issue, plenty of learning methods have been proposed over the years, ranging from new bio-inspired learning rules to adaptations of the gradient descent algorithm to deal with the aforementioned discontinuity, among others. To top it all off, each learning method comes with its own limitations in terms of input encoding scheme, neuron models, network architectures and network running mode, among others [27]. In this subsection we briefly revisit them as a referential knowledge base for concepts later referred to in the remainder of the article.

Encoding methods for SNNs can be divided into two big groups: rate encoding schemes and temporal encoding schemes [28]. In both cases, their purpose is to transform input samples into spike trains that can be processed by the network. As such, rate encoding schemes embeds input data information in the instantaneous or averaged rate of generated spikes in a single or group of neurons, i.e., in the count of spikes emitted within a time encoding window. On the other hand, temporal encoding methods conveys the input information in the precise timing of the spikes.

As for the neuron models, a great variety exists, each with its own biological plausibility and processing efficiency. Basically, they describe the evolution of the membrane potential of the neuron and the spike generation at different levels of detail. The Leaky Integrate-and-Fire (LIF) model, for instance, is the most implemented model in the SNN community due to its low computational cost and simplicity when describing neuron dynamics [29]. Basically, the neuron integrates the input spikes as currents applied to a resistor-capacitor (RC) circuit that increases its voltage until a threshold is exceeded. At this point in time, a spike is emitted and the voltage of the neuron is reset to the resting value. The LIF neuron abstracts away the shape and profile of the output spike: it is simply treated as a discrete event. As a result, information is not stored within the spike, but rather the timing (or frequency) of spikes. The Spike-Response Model (SRM) behaves similarly to the LIF model, with the main difference being that it has an adjustable threshold that depends on the time since the last post-synaptic spike occurred. While LIF models are usually defined in terms of differential equations, the SRM expresses the membrane potential at a certain time as an integral over the past [30]. More biologically plausible models include the Hodgking-Huxley or the Izhikevich models [31], among others.

As in traditional neural networks, SNNs can run in an offline or online fashion. SNNs have shown their ability to build spike-time learning rules that capture temporal associations by leveraging spike information representation. Therefore, not only they are well suited for learning from spatio-temporal data in a batch learning setting, but they are also inherently effective for handling these temporal associations in streaming data [32].

After the above short primer on concepts related to SNNs, we proceed by describing several learning algorithms relevant to our study. It is not the aim of this work to comprehensively review all existing learning methods for SNNs. Therefore, some few insights about the most relevant approaches for this purpose are next provided. To begin with, the Spike-Time Dependent Plasticity (STDP) [33] is a strongly

bio-inspired unsupervised learning rule that strengthens the connection between pre- and post-synaptic neurons that fire one after the other and that has inspired many other learning algorithms. SpikeProp [34] was the first to train SNNs by backpropagating errors, done on a single hidden layer architecture, with the SRM neuron model and the limitation that neurons were limited to emitting only one spike. Works thereafter refined this method to enable multiple spike firing at desirable times [35]. The main idea of the Spike Pattern Association Neuron (SPAN) [36] is to transform the input and the desired output spike trains into analog signals by convolving the spikes with a kernel function. In this manner, the computation of error signals is simplified, allowing for the application of gradient descent to optimize the synaptic weights.

However, learning methods discussed so far have difficulties to effectively train deep (multi-layered) SNNs. Surrogate Gradient (SG) methods overcome the discontinuity of the spiking neuron model by substituting it with a smooth continuous relaxation. Therefore, it does not impose any restrictions in the learning algorithm that can be used to train the model. Consequently, SG methods can be combined with Back-Propagation Through Time (BPTT) to give rise to recurrent flavors of SNNs. Nevertheless, this workaround poses stringent limitations for SNNs in terms of computational and memory resources when deployed over traditional hardware (GPUs). Moreover, its implementation over neuro-morphic hardware may impose locality requirements that can complicate its deployability [37]. A remarkable method that solves the latter and utilizes a surrogate gradient approach is *SuperSpike*, a biologically plausible nonlinear Hebbian three-factor rule with individual synaptic eligibility traces [38]. Basically, the objective of the algorithm is to minimize the van Rossum distance between the desired and the actual spike train through gradient descent. Hence, the derivatives of the spikes with respect to the weights of the network are needed. To make this computation feasible, the spike train is approximated by replacing it with a continuous auxiliary function that depends on the difference between the membrane potential of the LIF neuron and its firing threshold.

2.3. Explainability in spiking neural networks

As in traditional neural networks, SNNs are inherently opaque models that need to be complemented with explanations and justifications of their issued predictions, particularly in critical applications involving a non-expert audience. Currently it is widely acknowledged that SNNs are not applied in practical modeling problems as often as their non-spiking counterparts. However, their event-based working procedure is making SNNs gradually replace non-spiking neural architectures in problems and tasks driven by computational efficiency. This progressively increasing prevalence of spiking neural computation approaches makes it of utmost necessity to derive XAI methods specific for this family of models, over as many tasks as possible, so that the time to their acceptance and adoption in real-world settings is effectively narrowed.

Despite this noted relevance, efforts invested so far in this direction are scarce. To begin with, the work in [39] proposed a new knowledge encoding method that enables the attribution of the input features in the classification decision in a Multi-Class Synaptic Efficacy Function based leaky-integrate neuron (MC-SEFRON) classifier, a model without hidden layers based on the one presented in [40]. MC-SEFRON leverages a time-varying weight neuron model and trains it by using a modified STDP rule previously designed in preceding contributions. The result is a classifier with great accuracy when tested over the MNIST image classification task, and capable of eliciting explanations in the form of attribution maps for input images. The downside is that explanations are tightly bounded to the crafted model, and therefore are not generalizable to other spiking neuron architectures in more frequent use by the related community. Methods for rule extraction

from SNNs are presented in [41] and also in [42,43], where the NeuCube architecture [44] is used.

When searching for a more general explainability method, the contribution [45] exploited the fact that spikes with short inter-spike intervals highly contribute to the neural decision process. This observation gave rise to the Spike Activation Map (SAM) proposed in this study. For each timestep and neuron, a Neural Contribution Score (NCS) aggregates the Temporal Spike Contribution Score (TSCS) of all the pre-synaptic spikes previously processed by every neuron. These NCS are calculated at every time step for the feature maps of a convolutional neural network, and are summed pixel-wise across all channels, obtaining the aforementioned SAMs. Hence, one heatmap is generated by this technique for every time step and SNN layer.

When shifting the focus on time series data, the contribution in [46] combines three components to design a local feature-based explanation called Temporal Spike Attribution (TSA): (i) the influence of spike times, modeled by the NCS introduced above; (ii) the influence of the model's parameters, extracted from the weights of the connections between layers; and (iii) the influence of the classification decision, inferred from the softmax values. To the best of our knowledge, this threefold approach is the more recent of the short series of contributions dealing with explainability techniques specifically devised for SNNs.

2.4. Contribution over the state of the art

After examining the current status of knowledge of the research fields tackled in our work, we note that multi-layered SNNs are lately becoming a viable option to implement and test in practical settings, performing competitively against non-spiking neural counterparts with notably improved energy efficiency. For instance, spiking-based CNN models comprising several neural processing layers have been proposed recently for different image classification tasks [47–50]. This is the reason why the community is slowly expressing interest in addressing practical issues of these models when used in real-world environments, including the detection of OoD samples. However, as far as our revision of the literature can tell, there is no prior work in which OoD detection is studied for SNNs, nor have explanations been fabricated for modeling tasks besides classification. This work steps beyond the current state of the art by designing a specific method to detect OoD samples for SNNs, validating that not only it can be done by examining information produced during the training process of the SNN, but also that it attains better detection statistics than post-hoc methods from the literature that can be directly applied to these models. Further along this line, OoD predictions issued by the proposed detector can be combined with those produced by other OoD detection methods so as to yield an ensemble of OoD techniques that improve further their individual performance scores. This exposes that predictions produced by the proposed approach are complementary, coherent and diverse enough w.r.t. to other techniques for its fusion to enhance the overall detection figures. In addition, we propose an attribution method aimed to make the OoD detection process more understandable and trustworthy for the user, which spans further the narrow gamut of modeling tasks for which XAI methods have so far been developed in related studies.

3. Proposed approach

We now describe the method designed to detect OoD samples in SNNs, which we hereafter denote as *Spike Count Pattern* (SCP) based detector. Its workflow is graphically sketched in Fig. 2. First, a mathematical formulation of the network characteristics is posed in Section 3.1. Next, an in-depth explanation of the detector is provided in Section 3.2. Finally, Section 3.3 describes the steps to produce an attribution map based on the information generated during the OoD detection performed by the SCP-based detector.

3.1. Spiking neural network: mathematical notation

Our method targets the detection of OoD instances by exploiting the internals of SNNs working with rate-based encoding strategies. We assume a training dataset D_{tr} with samples $\{\mathbf{x}^{(i)}\}_{i=1}^M$, where $\mathbf{x}^{(i)} \in \mathbb{R}^D$, M denotes the number of samples in the training set, and D is the dimension of every input instance. Each of such training samples belong to a class $c \in \{1, 2, \dots, C\}$, where C denotes the number of classes in D_{tr} . The SNN model requires spike trains $\mathbf{s}^{(i)}(t) = \{s_d^{(i)}(t)\}_{d=1}^D$ at their input which, when implemented in discrete-time processing hardware, are represented by binary sequences $\mathbf{s}^{d,(i)} = \{s_1^{d,(i)}, s_2^{d,(i)}, \dots, s_{T/\Delta T}^{d,(i)}\}$, where $s_t^{d,(i)} \in \{0, 1\}$ and a one represents the occurrence of a spike. Each sample will be processed during a certain amount of time, named the simulation time T , that is divided in time steps of width ΔT . The number of time steps within a simulation period is easily obtained as $T/\Delta T$, which establishes the length of the binary sequence in which spikes are represented in the discrete domain.

Hence, inputs must be converted in spike trains by an encoding method $g(\mathbf{x}^{(i)}) : \mathbb{R}^D \mapsto \{0, 1\}^{t \times D}$, where the addition of the time dimension is done if the data is static. Here a rate-based coding scheme that converts each feature to a spike train using an homogeneous Poisson process is employed [51]. Basically, for each feature, the probability of a spike $s_t^{d,(i)}$ being emitted at time $t \in \{1, \dots, T/\Delta T\}$ for feature $d \in \{1, \dots, D\}$ and instance $\mathbf{x}^{(i)}$ is given by:

$$Pr\{s_t^{d,(i)}\} = r\Delta T = x_d^{(i)} r_{max} \Delta T, \quad (1)$$

where ΔT is the interval width, and r denotes the firing rate whose value depends on the value of each normalized feature $x_d^{(i)} \in \mathbb{R}[0, 1]$ and a maximum firing rate threshold r_{max} . The value of r_{max} is upperbounded by $1/\Delta T$ to avoid generating more than one spike per interval. To generate spike trains, at each timestep a sample $z_t^{d,(i)}$ is drawn uniformly at random from a continuous distribution $\mathcal{U}(0, 1)$, and is compared to the probability in Expression (1), such that a spike is generated if $z_t^{d,(i)} \leq r\Delta T$. As a result, the obtained spike train $\mathbf{s}^{(i)}(t)$ comprises randomly distributed spikes, such that the number of spikes in the simulation period T follows a Poisson distribution.

To process an instance $\mathbf{x}^{(i)}$, the SNN converts each of its features $x_d^{(i)}$ into a spike train $\mathbf{s}^{d,(i)}$. At each time step t , the corresponding part of all the spike trains $\{s_1^{1,(i)}, \dots, s_1^{D,(i)}\}$ are fed to the model. Assuming a Leaky-Integrate-and-Fire neuron model, the state of the LIF neurons of each SNN layer is updated depending on the occurrence of a spike in the previous layer and the previous state of the neuron. If an activation threshold is surpassed, the potential/voltage of the neuron is reset to its resting state value for the next simulation time step $t+1$. It is important to note that spikes affect the next layer's neuron states through the so-called *synaptic current*, which impacts on the membrane voltage of the neuron. This synaptic current generated by a spike depends only on the weight of the connection between the actual neuron (post-synaptic) and the neurons of the previous layer (pre-synaptic).

When the SNN is used for classification tasks, the output layer is not composed by LIF neurons, but rather by leaky integrators matching the number of classes in D_{tr} . This type of neurons only accumulates voltage through time, i.e., it lacks any mechanisms to fire a spike. The maximum voltage levels achieved by these neurons during the simulation period T are interpreted as the logits input to a softmax function, from where the corresponding predicted class is inferred under a *winner-takes-all* decision strategy (as in classical non-spiking neural classifiers).

3.2. SCP-based Out-of-Distribution detection approach

The design of the proposed OoD detection approach departs from the intuitive observation that similar samples of the same class tend to have similar neuron activation patterns in the hidden layers. Indeed, the SCP-based detector is largely inspired by this observation to yield a novel OoD detection method, which is graphically depicted in Fig. 3. In

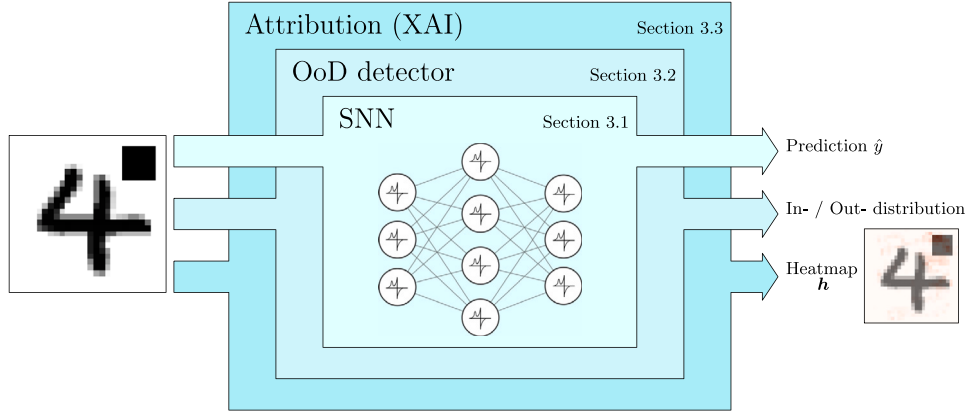


Fig. 2. Workflow of the SCP based detector for an image classification task.

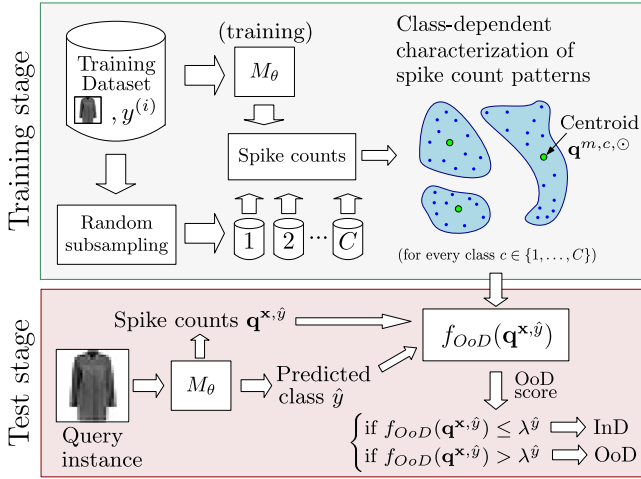


Fig. 3. Graphical description of the methodology of the SCP based detector.

essence, the detection hinges on comparing the spike count pattern of every new test sample to the representative (archetypical) spike count patterns of its predicted label, which can be characterized during the training stage of the SNN. Therefore, the method requires augmenting the training algorithm of the SNN in order to characterize and produce such spike count archetypes, so that they can be used at inference time to detect whether any test instance is out of the training distribution. Here, a spike count $q(s)$ corresponding to a spike train $s \in \{0, 1\}^{T/\Delta T}$ is given by the number of spikes fired during the simulation period, namely, $q(s) = \sum_{t=1}^T s_t$, where we recall that $s_t \in \{0, 1\}$ and value 1 represents a fired spike.

In doing so, and for the sake of a reduced computational complexity added to the training process, the SCP-based detector selects uniformly at random a specific number of training examples of every class and collects the spike counts of the layer prior to the softmax (for reasons disclosed later). Once spike counts have been computed, a class-conditional clustering is performed over the spike counts, using a distance metric to gauge the dissimilarity between spike counts. After computing all pairwise distances, the centroid of each class-conditional cluster resulting from the clustering stage is computed, which can be regarded as the aforementioned spike count archetypes of the samples inside the cluster of the class to which they belong. It should be noted that every class in D_{tr} may span more than one cluster, depending on the inherent intra-class variability of the spike counts produced by the training instances of the class at hand. Finally, at inference time, a test instance x' is processed through the already trained SNN, yielding a predicted class \hat{y} . A score is then computed to determine whether a

new sample belongs to the in- or out-distribution, which is given by the closest L_1 distance to the archetypes of its predicted class \hat{y} . If this distance is greater than a predefined threshold λ , x' is declared to be an OoD sample. Otherwise, the SCP-based detector determines that x' belongs to the distribution of the training data.

Algorithm 1: Proposed SCP detector

Data: Training data D_{tr} , size P of training samples D_{tr}^c , last layer index L , spikes $s^{n,L(i)}$ associated to example $x^{(i)}$, distance measure $d(\cdot)$, clustering algorithm, aggregation function $f_{agg}(\cdot)$, OoD score function $f_{OoD}(\cdot)$, thresholds $\lambda^{\hat{y}}$, query instance x

Result: Decision on the OoD nature of x

- 1 //Training stage
 - 2 Train SNN over D_{tr}
 - 3 **for** $c \in \{1, \dots, C\}$ **do**
 - 4 $D_{tr}^c \leftarrow P$ sized c -class sample randomly selected from D_{tr}
 - 5 Collect the spike trains $s^{n,L(i)} = \{s_t^{n,L(i)}\}_{t=1}^{T/\Delta T}$
 - 6 Compute the spike counts $q(n, L, (i), c)$ as per Expression (2)
 - 7 Aggregate spike counts into a vector
 - 8 $\mathbf{q}^{(i),c} = \{q(n, L, (i), c)\}_{n=1}^{N_L}$
 - 9 Compute distances $d(\mathbf{q}^{(i),c}, \mathbf{q}^{(j),c}) \forall \mathbf{q}^{(i),c}, \mathbf{q}^{(j),c} \in D_{tr}^c$
 - 10 Extract M^c clusters $\{\mathcal{Q}^{m,c}\}_{m=1}^{M^c}$, where $\mathcal{Q}^{m,c} \subseteq \{\mathbf{q}^{(i),c}\}_{i \in D_{tr}^c}$
 - 11 **for** $m \in \{1, \dots, M^c\}$ **do**
 - 12 | Compute archetype $\mathbf{q}^{m,c,0}$ applying $f_{agg}()$ to $\mathcal{Q}^{m,c}$;
 - 13 //Detection stage
 - 14 Let $\hat{y} = M_{\theta}(x)$ (prediction of query sample)
 - 15 Collect spike trains $s^{n,L,x}$
 - 16 Compute spike count $\mathbf{q}^{x,\hat{y}} = \{q(n, L, x, \hat{y})\}_{n=1}^{N_L}$
 - 17 Assign an OoD score by computing $f_{OoD}(\mathbf{q}^{x,\hat{y}})$ as per Expression (4)
 - 18 Declare x to be OoD if $f_{OoD}(\mathbf{q}^{x,\hat{y}}) > \lambda^{\hat{y}}$, otherwise x follows the in-distribution of D_{tr}
-

We proceed by describing step by step the SCP-based detector. As schematically described in Algorithm 1, our developed technique requires a training stage for generating archetypical spike count patterns for every class in the dataset. For simplicity we assume L stacked layers in the SNN architecture, comprising N_L LIF neurons. By slightly modifying the previous notation, we denote the time slot of the spike train occurring at the n th neuron of the last layer L (the layer prior to the leaky integrator) as $s^{n,L(i)} = \{s_t^{n,L(i)}\}_{t=1}^{T/\Delta T}$. Based on this definition, the SCP-based detector starts by extracting P spike trains $\{s^{1,L(i)}\}_{i \in D_{tr}^c}$ for each of the labels in the training dataset (line 5 of Algorithm 1). To this end, a P -sized subset $D_{tr}^c \subset D_{tr}$ of the training instances predicted

to belong to label c is sampled at random (line 4), such that $|D_{ir}^c| = P \forall c \in \{1, \dots, C\}$. Once such spike trains are collected, counts can be computed for each spike train at neuron n and instance $\mathbf{x}^{(i)}$ (line 6) as:

$$q(n, L, (i), c) = \sum_{t=1}^{T/\Delta T} s_t^{n,L,(i)}, \text{ for } i \in D_{ir}^c, \quad (2)$$

which can be aggregated into a spike count vector $\mathbf{q}^{(i),c} = \{q(n, L, (i), c)\}_{n=1}^{N_L}$ (line 7). The proposed SCP-based detector extracts the spikes of the *last layer* L of the SNN architecture for two reasons. The first is that the deeper the selected layer is along the hierarchy of neural layers, the higher the level of the features handled is, and the lower the number of spike count trains will be. High-level semantics are better for OoD sample discriminability, whereas low-dimensional spaces are preferred for characterizing the space of archetypical behaviors of the activations of instances of a given class. Second is the fact that the softmax layer is bounded to represent the probability of the input belonging to a certain class, which may not be well suited to characterize OoD samples.

Based on the P spike count vectors $\{\mathbf{q}^{(i),c}\}_{i \in D_{ir}^c}$ computed for every class c as per Expression (2), a clustering algorithm can be used over such spike counts to group them in terms of their inter-vector similarity (line 9). This clustering process yields M^c clusters $\{\mathcal{Q}^{m,c}\}_{m=1}^{M^c}$, where $\mathcal{Q}^{m,c} \subseteq \{\mathbf{q}^{(i),c}\}_{i \in D_{ir}^c}$ denotes the m th cluster arising from the spike count activations of the random subsample drawn from the training instances predicted to belong to class c . To compute the dissimilarity between spike count vectors we select the L_1 norm (also referred to as *Manhattan distance*) between two vectors (line 8), namely:

$$d(\mathbf{q}^{(i),c}, \mathbf{q}^{(j),c}) = \sum_{n=1}^{N_L} |q(n, L, (i), c) - q(n, L, (j), c)|, \quad (3)$$

as its behavior in high-dimensional spaces is known to behave better than other alternatives, e.g., the Euclidean norm [52]. Finally, all samples in every cluster $\mathcal{Q}^{m,c}$ are aggregated to obtain an archetype $\mathbf{q}^{m,c,\odot}$ (line 11), which is considered the representative spike count pattern of all samples in the corresponding cluster. We do so by devising a function $\mathbf{q}^{m,c,\odot} = f_{agg}(\mathcal{Q}^{m,c})$, which can be chosen to be any aggregation strategy that allows for a variable-size set of spike count patterns in their argument, e.g. the mean, the median or any other statistic alike.

After computing the archetypes of every cluster and class $\mathbf{q}^{m,c,\odot}$, the OoD detection process is ready to process each new query or test sample \mathbf{x} fed to the model. First, its prediction \hat{y} is output by the trained SNN model (line 13). Along the inference process, spike trains $\mathbf{s}^{n,L,\mathbf{x}}$ are collected at the last layer of the network (line 14), from where spike counts vectors $\mathbf{q}^{\mathbf{x},\hat{y}} = \{q(n, L, \mathbf{x}, \hat{y})\}_{n=1}^{N_L}$ are computed (line 15). Next, an score is assigned to the test instance by using a scoring function (line 16):

$$f_{OoD}(\mathbf{q}^{\mathbf{x},\hat{y}}) = \min_{m \in \{1, \dots, M^{\hat{y}}\}} d(\mathbf{q}^{\mathbf{x},\hat{y}}, \mathbf{q}^{m,\hat{y},\odot}) \quad (4)$$

namely, the minimum L_1 distance between the spike count pattern $\mathbf{q}^{\mathbf{x},\hat{y}}$ of test instance \mathbf{x} and the archetypes $\mathbf{q}^{m,\hat{y},\odot}$ for $m \in \{1, \dots, M^{\hat{y}}\}$ computed for its predicted class \hat{y} . Finally, the value of $f_{OoD}(\mathbf{q}^{\mathbf{x},\hat{y}})$ is used to decide whether instance \mathbf{x} belongs to the in- or the out-distribution (line 17). If the score is greater than a certain class-conditional threshold $\lambda^{\hat{y}}$, the sample is considered as an out-of-distribution sample, declaring the sample from the in-distribution otherwise:

$$\mathbf{x} \text{ belongs to } \begin{cases} \text{in-distribution} & \text{if } f_{OoD}(\mathbf{q}^{\mathbf{x},\hat{y}}) \leq \lambda^{\hat{y}}, \\ \text{out-distribution} & \text{if } f_{OoD}(\mathbf{q}^{\mathbf{x},\hat{y}}) > \lambda^{\hat{y}}. \end{cases} \quad (5)$$

3.3. Local relevance attribution method

In case the query instance \mathbf{x} is declared not to belong to the in-distribution, the SCP-based detector is complemented by a local attribution-based explanation technique. This explanatory technique is local, as it operates over a given test instance towards generating a

relevance attribution vector $\mathbf{h}(\mathbf{x}) \in \mathbb{R}^D$, where each component $h_d(\mathbf{x})$ quantifies the importance of each feature of the input when detecting it as an OoD sample.

The steps of this local attribution method are schematically illustrated in Fig. 4 and detailed in Algorithms 2 and 3. In the following, we refer as *input layer* to the first fully-connected layer of the SNN network. In the case a feature extractor (e.g. a convolutional layer) is preceding that fully connected stage, the input layer will refer to the output of the feature extractor.

The core ideas driving the design of this method are two: on one hand, we aim to reproduce the detector's procedure to detect OoD samples in the input space of the SNN rather than in the latent space where the detection is done (namely, the L th layer prior to the softmax); on the other hand, since an attribution is sought, the method must account for the influence of each input feature on the aforementioned latent space. The latter is achieved by *propagating* the spike counts used for OoD detection at layer L back to layer 1 by using the weight matrices $\mathbf{W}^{[l]} : l \in [1, L-1]$ of the network; and the former is achieved using a L_1 distance over the backpropagated spike counts. Here, l is the layer index.

The local attribution mechanism proceeds as follows: first, we back-propagate the spike counts of both the input $\mathbf{q}^{\mathbf{x},\hat{y}}$ and the centroid $\mathbf{q}^{m,\hat{y},\odot}$ to the input layer using the weight matrices $\mathbf{W}^{[l]}$. This is accomplished by recurrently applying a function f_{bp} defined as:

$$\mathbf{q}^{\mathbf{x},\hat{y},[l-1]} = f_{bp}(\mathbf{W}^{[l]}, \mathbf{q}^{\mathbf{x},\hat{y},[l]}) = \mathbf{W}^{[l]T} \cdot \mathbf{q}^{\mathbf{x},\hat{y},[l]}, \quad (6)$$

where we extend previous notation¹ with superindex $[l]$ to indicate that spike counts correspond to layer $l \in \{1, \dots, L-1\}$. This function can be viewed as the mapping of the spike counts of one layer to the space of the previous layer. Specifically, the spike counts are multiplied by the transpose $\mathbf{W}^{[l]T}$ of the weight matrix at layer l . This operation simultaneously transforms the spike counts at the output of the layer back to its inputs, and implicitly implements a branching process similarly to what occurs in the backward pass of gradient calculus over computational graphs. Therefore, recurrently applying this function over all layers between the penultimate layer ($L-1$) and the input layer ($l=1$), we end up with the spike counts mapped to the input space, obtaining what we have called the *reconstruction* of the spike counts in the input space. This operation is described in Algorithm 2, which performs this mapping.

Algorithm 2: Reconstruction of spike counts at the input layer

Data: Weight matrices $\mathbf{W}^{[l]}$, spike counts $\mathbf{q}^{\mathbf{x},\hat{y},[L]}$, function to map spike counts to previous layer $f_{bp}(\cdot)$

Result: Spike counts mapped to input space $\mathbf{q}_{rec}^{\mathbf{x},[1]} \in \mathbb{R}^D$

- 1 $\mathbf{q}_{rec}^{\mathbf{x},[L-1]} = f_{bp}(\mathbf{W}^{[L]}, \mathbf{q}^{\mathbf{x},\hat{y},[L]})$
 - 2 **for** $l \in \{L-2, \dots, 1\}$ **do**
 - 3 $\mathbf{q}_{rec}^{\mathbf{x},[l]} = f_{bp}(\mathbf{W}^{[l+1]}, \mathbf{q}_{rec}^{\mathbf{x},[l+1]})$
 - 4 **Let** $\mathbf{q}_{rec}^{\mathbf{x},[1]} = \{\mathbf{q}_{rec,d}^{\mathbf{x},[1]}\}_{d=1}^D$
 - 5 **for** $d \in \{1, \dots, D\}$ **do**
 - 6 $\mathbf{q}_{rec,d}^{\mathbf{x},[1]} = \max(\mathbf{q}_{rec,d}^{\mathbf{x},[1]}, 0)$
-

It is important to remark the reason behind the maximum or ReLU operation done in the last step of Algorithm 2. To obtain a reasonable attribution, our goal is to reproduce the operating conditions of our detector, but in the input space rather than in the latent space. As LIF neurons are in use (whose output can be either a spike or nothing), when we integrate them in time (computing the spike counts) we will always obtain a positive or a 0 value, even if the weights arriving a neuron are all negative i.e. inhibitory. Therefore, to imitate this behavior when backpropagating static values (counts), negative values obtained from Expression (6) should be set to zero.

¹ Accordingly, the spike counts $\mathbf{q}^{\mathbf{x},\hat{y}}$ over which OoD detection is done as per Algorithm 1 should be denoted as $\mathbf{q}^{\mathbf{x},\hat{y},[L]}$.

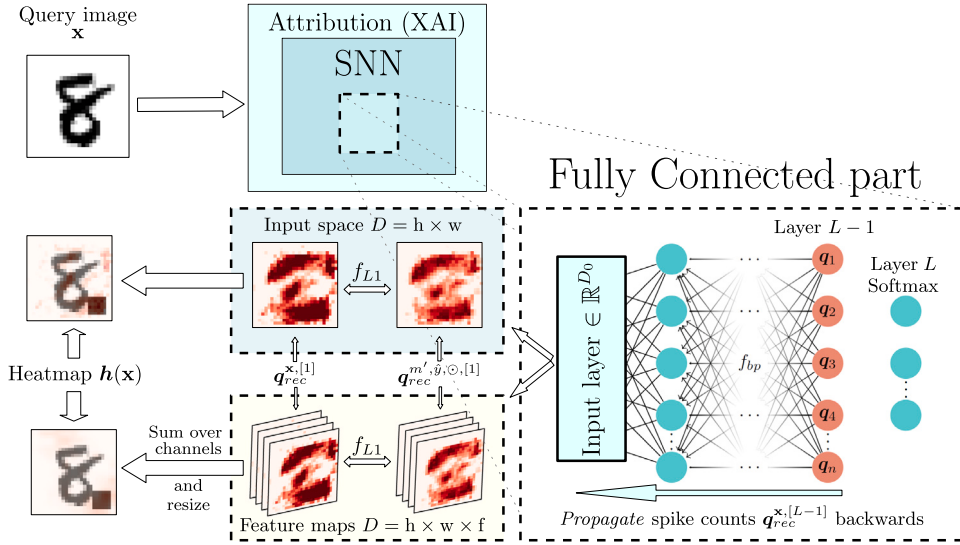


Fig. 4. Schematic description of the local attribution method. Here, h , w and f refer to height, width and filters respectively.

Algorithm 3: Construction of the relevance attribution vector

Data: Weight matrices $W^{[l]}$, spike counts $q^{x,\hat{y},[L]}$ and the closest centroid $q^{m',\hat{y},\odot}$ among $\{q^{m,\hat{y},\odot}\}_{m=1}^{M^{\hat{y}}}$ resulting from Algorithm 2

Result: Relevance attribution vector $h(x) \in \mathbb{R}^D$

- 1 Use Algorithm 2 to obtain s ;
- 2 Repeat Algorithm 2 to extract the reconstructed centroid $q_{rec}^{m',\hat{y},\odot,[1]}$ from $Q^{m',\hat{y}}$ for the closest centroid $m' \in \{1, \dots, M^{\hat{y}}\}$ to $q^{x,\hat{y},[L]}$;
- 3 **for** $d \in \{1, \dots, D\}$ **do**
- 4 $h_d(x) = |q_{rec,d}^{x,[1]} - q_{rec,d}^{m',\hat{y},\odot,[1]}|$;

Second, after obtaining the reconstructions $q_{rec}^{x,[1]}$ and $q_{rec}^{m',\hat{y},\odot,[1]}$ (where m' stands for the closest centroid to the spike counts corresponding to sample x), we seek how to reproduce the OoD detector's operation in this input space. To this end we resort to a L_1 distance as in Expression (3) so as to compute the absolute difference of each feature of both reconstructions, namely:

$$h_d(x) = |q_{rec,d}^{x,[1]} - q_{rec,d}^{m',\hat{y},\odot,[1]}|, \quad \forall d \in \{1, \dots, D\}, \quad (7)$$

i.e., the relevance attribution vector results from the application of the same L_1 -based OoD detection criterion used in the proposed SCP-based approach, yet in the reconstruction of the spike counts in the input space.

To end with the description of the proposed relevance attribution technique, it is important to note that, depending on how the input data is encoded to the hierarchy of fully-connected LIF neural layers, an additional step would be needed to map $h(x)$ (namely, a D -dimensional vector) to the dimensions of the original data. We exemplify this additional operation by describing two cases:

- Let us assume that x is an image, so that it is input to the network by means of a fully-connected input layer with as many LIF neurons as the number of pixels and channels of the image. In this case, $h(x)$ should be *unflattened* in the same order as the one followed to serialize the input image into the network.
- If a convolutional layer is used to extract features from the input image, the dimension D of the relevance attribution network depends on the size of the convolutional filters and their cardinality. In this case, the rearrangement of the relevance attribution vector operates

first by rearranging its components to the shape of the feature maps of the convolutional layer, yielding a three-dimensional tensor-like structure with width and height equal to the size of the convolutional filters, and with depth equal to the number of convolutional filters. Then, we perform a depth-wise aggregation of this tensor. Finally, the obtained map is scaled up by interpolation to generate a relevance attribution heatmap with the size of the original input image.

4. Experimental setup

In order to examine the performance of the proposed detection and attribution techniques, we design a comprehensive experimental setup using SNNs for image classification, aimed to inform with empirical evidence the responses to the following research questions:

- RQ1: Does the proposed SCP-based detector perform competitively when compared to post-hoc OoD detectors from the literature adapted to SNN architectures?
- RQ2: How do the SCP-based detector perform when fused with other adapted OoD detection methods?
- RQ3: Does the local relevance attribution method yield informative explanations about the features that make a given sample be detected as OoD? Which are the limitations of this explainability technique?

The design of the experiments devised to answer the above questions include the definition of the SNN architectures (Section 4.1), the in-distribution and out-distribution datasets under consideration (Section 4.2), the OoD methods selected for comparison and the scores used to measure their performance (Section 4.3), and the methodology followed to qualitatively validate the relevance attribution vectors generated for OoD samples by the proposed SCP-detector (Section 4.4). In what follows, details about these design choices are given. All scripts and notebooks producing the results reported in this manuscript have been made publicly available in a GitHub repository (https://github.com/aitor-martinez-seras/OoD_on_SNNs).

4.1. Considered SNN architectures

We start describing the experimental setup by the specific implementation of the SNN under consideration. SNNs comprise a very wide research field that lies in between Neuroscience and ML. Hence, some SNN variants are designed departing from the bio-plausibility perspective, attempting to make a machine simulate faithfully the biological brain, neurons and its synapses. Other variants, however, pursue

Table 1

SNN architectures considered in the experiments for gray-scale datasets. Each layer is defined by a number and some letters, which respectively define the number of neurons and the type of connections. Each layer is separated by a hyphen. E denotes Poisson encoding neurons, fc stands for fully connected neurons, conv means convolutional filters and Flt refers to a flatten. The subscript of the last two defines the voltage threshold ϑ_{th} of the LIF neurons. In the last layer only leaky integrators (LI) are used, followed by a SoftMax activation, being the number of output neurons equal to the number of classes in the in-distribution dataset.

Network type	Reference	Architectures
Fully Connected SNN (FC-SNN)	FC_1	784E – 200fc _{0.25} – LI
	FC_2	784E – 300fc _{0.25} – 200fc _{0.25} – LI
Convolutional SNN (CNN-SNN)	CNN_1	784E – 20conv _{0.2} avgpool – 50conv _{0.2} – Flt – 300fc _{0.1} – LI
	CNN_2	784E – 20conv _{0.2} avgpool – 50conv _{0.2} – Flt – 500fc _{0.1} – 300fc _{0.05} – LI

SNNs capable of solving complex modeling tasks due to their inherent advantages rather than the biological plausibility of the dynamics and neurons used. The latter design criterion has dominated this area in recent times due to the need for more efficient ML algorithms. As a result, a great variety of SNN frameworks are nowadays available [53]. Although some of them seek a realistic simulation of the human brain, others offer simpler albeit more effective SNN implementations for data-based modeling, in terms of the depth of the networks that can be trained and the complexity of the tasks they can tackle.

As stated in the introduction, our experiments focus on practical fully-spiking SNN implementations, namely, networks that are trained and run entirely on data encoded as temporal spikes. Specifically, we use 2nd order LIF neurons that account for synaptic conductance. By denoting the membrane voltage as ϑ and the synaptic current as ζ , these decay over time at rates τ_{mem} and τ_{syn} , respectively. The ordinary differential equations that are integrated over time (in discrete time steps) are [54]:

$$\dot{\vartheta} = 1/\tau_{mem}(\vartheta_{leak} - \vartheta + \zeta), \quad (8)$$

$$\dot{\zeta} = -1/\tau_{syn}\zeta, \quad (9)$$

together with the jump condition:

$$z = \Theta(\vartheta - \vartheta_{th}), \quad (10)$$

and transition equations:

$$\vartheta = (1 - z)\vartheta + z \cdot \vartheta_{reset}, \quad (11)$$

$$\zeta = \zeta + \zeta_{in}, \quad (12)$$

where ϑ_{th} defines the membrane voltage that, when exceeded, leads to the emission of a spike; z is a variable taking value 1 if the neuron voltage surpasses ϑ_{th} (0 otherwise); ϑ_{leak} refers to a constant voltage value that is leaked from the neuron; $\Theta(\cdot)$ stands for the Heaviside step function that returns value 1 when the threshold is surpassed; and ϑ_{reset} is the value of voltage that is set when the latter happens. A SNN framework that features this LIF model is Norse [55], which runs on top of PyTorch and provides rate-based encoding strategies, LIF-based and Izhikevich neuron models, and the SuperSpike training algorithm mentioned in Section 2.2.

We have tested our method over two different types of SNN architectures for image classification: a Fully Connected (FC) architecture and a Convolutional Neural Network (CNN) architecture. In both cases, the encoder used to translate pixel values to spike trains is the so-called Poisson encoder, whereas neuron models are LIF (intermediate neurons) except for the output layer, which utilizes Leaky Integrators (LI) followed by SoftMax. LI cells are similar to LIF neurons, but never fire: they only accumulate voltage. Parameters for both types of cells are $\vartheta_{leak} = 0$, $\vartheta_{reset} = 0$, $\tau_{mem} = 0.005 \text{ ms}^{-1}$, and $\tau_{syn} = 0.01 \text{ ms}^{-1}$. These parameters are identical for all cells in the SNN, except for the voltage threshold ϑ_{th} , which is specified layer-wise in Tables 1 (gray-scale datasets) and 2 (color datasets) along with the architectures of the networks in use for every case. Information in these tables can be interpreted as explained in their captions.

All gray-scale models are trained for 5 epochs using an Adam solver with learning rate equal to 0.002. The simulation time T is 50 ms,

whereas the time step duration ΔT is 1 ms, hence yielding 50 time steps per sample. For the model used over color datasets, the simulation time is expanded to 64 time steps per sample, whereas AdamW optimizer was used for a total of 75 epochs.

4.2. In-distribution and out-distribution datasets

In order to evaluate the performance of the proposed SCP-based detector, we have followed common practice in the literature related to OoD detection. Consequently, datasets of well-known image classification tasks have been considered as in-distribution datasets. Specifically, for the gray-scale datasets we have used 28×28 centered images: the MNIST dataset [56], composed by 60,000 images of handwritten images with 10 classes (digits from 0 to 9); Fashion MNIST [57], which comprises 60,000 images of 10 different classes of clothes; KMNIST [58], consisting of 60,000 images of Japanese characters (*Kanjis*), each belonging to one among 10 classes; and finally, EMNIST Letters [59], which provides 145,000 images of 26 types of letters of the Latin alphabet. In the case of color images, CIFAR10 and SVHN are used, composed by 32×32 images and 10 classes each, the former comprising animals and vehicle classes and the latter house numbers from 0 to 9 obtained from Google Street View images. In addition, the techniques are evaluated over neuromorphic image classification datasets, i.e., datasets whose examples are composed by event-based data coming, for example, from a DVS. The selected in-distribution datasets are NMNIST [60], a spiking version of the original frame-based MNIST dataset; and DVSGesture [61], a dataset created for developing a gesture recognition system with 11 hand gestures from 29 subjects under 3 illumination conditions.

When it comes to out-of-distribution datasets, we consider all possible combinations of pairs of the in-distribution datasets defined above, establishing one of them as the dataset over which the SNN is trained, and querying the trained SNN with the others, independently done for the gray-scale, color and neuromorphic datasets. For the sake of a wider experimental coverage, we have added several additional out-of-distribution datasets for both gray-scale and RGB cases. For the first case, we have considered notMNIST [62], which gathers font glyphs extracted for the letters A through J; omniglot [63], which amounts to 1,623 different handwritten characters from 50 different alphabets; and CIFAR10-BW, the black-and-white rescaled version (to 28×28 pixels) of the CIFAR10 dataset [64]. As a result, the SNN trained for every in-distribution dataset will be queried with images drawn from 6 different out-of-distribution datasets. In all cases, the OoD dataset is reduced to a uniform stratified sample of 10,000 samples. For the case of colored datasets, we add Food101, consisting of 101 food categories; LSUN (crop) and LSUN (resize), 2 small subsets of the LSUN scene classification dataset, the former composed by randomly cropping image patches of size 32×32 and the latter by downsampling each image to size 32×32 ; Flowers102, 102 categories of commonly occurring plants of United Kingdom; Caltech101, a collection of pictures of objects belonging to 101 classes; GTSRB, which comprises 43 classes of German traffic signs; TinyImageNet, a smaller version of ImageNet with 200 classes and images downsampled to 64×64 pixels; EuroSAT, a land cover classification dataset using Sentinel-2

Table 2

SNN architectures considered in the experiments for color datasets, following the same structure and legend as in Table 1. The new keyword **cpa** refers to a Convolutional layer with padding equal to 1 and average pooling.

Network type	Reference	Architectures
Convolutional SNN (CNN-SNN)	CNN_{rgb}	$784E - 32cpa_{0.25} - 64cpa_{0.2} - 128cpa_{0.1} - Flt - 2048fc_{0.1} - 100fc_{0.1} - LI$

satellite images; and finally, DTD, which is composed by images of 47 classes of different textures. Finally, we consider as OoD neuromorphic datasets CIFAR10DVS [65], the neuromorphic version of CIFAR10; and DVSLip [66], an event-based lip-reading dataset.

4.3. Methods for comparison and performance scores

For comparison purposes, to the best of our knowledge there is unfortunately no OoD detector specifically designed for SNNs. Consequently, our benchmark considers post-hoc OoD techniques that essentially rely on the output of the neural network being monitored, either the logits or the output of the softmax activation². In the case of the SNN architectures selected for our experimentation, logits correspond to the maximum voltage reach by every LI cell during the simulation time. The methods against which the SCP-based detector is compared include the Baseline method [14], ODIN [15] and the energy-based approach presented in [19]. The Baseline method does not require any parameter tuning, whereas ODIN and the energy-based method are governed by several parameters such as the temperature or the input preprocessing (in the case of ODIN). Values of these parameters must be properly set in order to obtain good results. To this end, for the energy-based detector we follow the study in [19] that states that can be used parameter-free by defining $T = 1$. For the temperature parameter of ODIN, as there is not such option, we have instead evaluated several temperature values for each (in-distribution, out-of-distribution) dataset combination, reporting on the configuration that yields the best performance in each case. For the input preprocessing, no perturbation is used, as its influence is minor when compared with temperature (as derived in [15]) and it depends on gradients which, for a fully-spiking SNN, can only be approximated.

Regarding the configuration of the proposed SCP-based detector, the subset of instances of each class used for the characterization of the spike count patterns is formed by $P = |D_{ir}^c| = 1000$ samples per class. Agglomerative hierarchical clustering was used to extract clusters from the spike counts of all training instances used for characterization. As for the creation of the centroids, a median statistic was selected as the aggregation function $f_{agg}(\cdot)$, which is robust against spike count outliers. Finally, the threshold for each class λ^c is selected utilizing a subset of 1000 samples per class extracted from the training data and different from the ones used for the creation of the centroids. To do so, the OoD scores are computed and the threshold λ^c is defined as the value that yields a target percentage of that class-wise subset that will be correctly classified as in-distribution.

To quantify the capability of the counterparts in the benchmark to detect OoD samples, several scores used in the literature related to OoD detection are measured:

- FPR95 (False Positive Rate at 95% True Positive Rate), that is, the rate of OoD samples wrongly classified as positive (in-distribution) when the True Positive Rate (TPR) is 95%. Here, TPR and FPR are defined in the usual way for binary classification, i.e., $TPR = TP/(TP+FN)$ and $FPR = FP/(FP+TN)$. In what follows a TP refers to an in-distribution sample correctly classified as such, whereas a TN stands for an OoD sample detected correctly by the OoD detector at hand.

² Other more recent approaches are not included because their implementation is not straightforward for fully-spiking architectures (Rankfeat [24]), nor have offline tests with these methods reported good detection scores that compensate for their high computational complexity (GradNorm [22]).




Fig. 5. Datasets for the experiments related to the quality of relevance attribution heatmaps.

- AUROC (Area Under the Receiver Operation Characteristic curve), which is a threshold-independent score for binary classification that can be regarded as the probability that the model ranks a random positive example with higher score than a random negative example. It is defined as TPR/FPR .
- AUPR (Area Under the Precision-Recall curve), another threshold-invariant score closely related to AUROC. The precision-recall curve is given by the relationship between $TP/(TP+FP)$ (precision) and TPR (recall).

4.4. Evaluating the quality of the relevance attribution method

When aiming to inform an answer to RQ3, there is no standardized way to measure the quality of explanations provided by a relevance attribution technique like the one proposed in this manuscript. To overcome this, we opt for evaluating our method against out-distribution datasets whose instances are built from in-distribution samples with imprinted visual artifacts. Since such artifacts are a priori known to be the parts of the image that *push most* the input sample towards the out-distribution, the relevance attribution technique should highlight such artifacts in the produced heatmap for the image at hand. In other words, if the difference between two samples is a certain extra and/or missing feature, our relevance attribution method should focus on them when performing OoD detection, and therefore should be clearly distinguishable in the heatmap.

In doing so, MNIST will be used as in-distribution dataset, whereas out-of-distribution datasets used for the qualitative attribution assessment include (i) MNIST-Square, that is, the MNIST dataset with a square of 5×5 pixels placed randomly in one of the corners of the image (with a separation to the border of the image of 2 pixels); and (ii) MNIST-C [67], another MNIST-based dataset that includes several modifications made to the digits. Specifically, we use the zig-zag corruption pattern in our experiments. An example showing the three datasets used in our experiments related to RQ3 is depicted in Fig. 5.

To correctly visualize the heatmaps, an appropriate plot range must be selected previously. To that end, a random sample is extracted from the training set and its heatmaps are computed. Then, from the pixel values of all the heatmaps, the maximum value and the q th quantile value are obtained, where q is the TPR selected to define the threshold. The maximum of the visualization range will be the calculated maximum value and the minimum will be the q th quantile value. The size of the random sample used to compute these maximum and quantile values is set to 100 examples.

Table 3

OoD detection scores obtained for the gray-scale FC-SNN architecture. The arrow next to every score indicates whether higher (↑) or lower (↓) values correspond to a better detection performance.

ID dataset	OoD dataset	AUROC ↑				AUPR ↑				FPR95 ↓			
		SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy
MNIST (96% acc)	FMNIST	92.91	74.19	74.19	39.71	91.41	65.67	65.67	41.12	23.15	76.34	76.34	93.00
	KMNIST	95.26	80.97	80.97	60.83	95.47	76.68	76.68	55.30	21.84	72.40	72.40	89.06
	Letters	88.38	84.25	84.25	69.17	88.83	83.31	83.31	65.79	50.33	71.15	71.15	89.40
	notMNIST	99.29	64.42	64.42	20.67	99.25	57.59	57.59	34.39	02.59	86.29	86.29	97.62
	omniglot	72.91	93.61	95.89	95.80	75.06	94.07	95.98	95.91	88.54	44.10	24.81	24.90
	CIFAR10-BW	99.60	80.13	80.13	16.92	99.58	78.94	78.94	33.34	00.96	80.51	80.51	99.68
FMNIST (84% acc)	MNIST	92.73	86.77	96.02	97.34	93.72	85.64	95.77	97.66	36.62	46.19	17.18	15.90
	KMNIST	86.23	78.79	88.37	87.29	85.90	81.75	88.76	87.73	50.72	79.00	57.01	65.11
	Letters	92.64	81.27	92.34	93.15	93.37	80.06	91.80	93.92	34.79	62.97	34.68	44.99
	notMNIST	98.53	61.80	61.80	43.30	98.97	62.28	62.28	44.74	02.67	93.22	93.22	95.99
	omniglot	79.25	92.16	98.07	98.70	81.49	91.61	98.16	98.98	78.62	30.63	08.01	06.00
	CIFAR10-BW	97.10	55.22	55.22	33.38	97.64	60.05	60.05	40.33	12.62	94.70	94.70	98.99
KMNIIST (87% acc)	MNIST	63.24	84.27	85.68	82.44	63.51	87.37	88.50	85.76	86.84	57.08	49.65	63.59
	FMNIST	72.11	76.58	76.58	46.86	67.25	78.31	78.31	49.50	63.47	70.38	70.38	90.74
	Letters	75.64	85.16	86.51	83.09	73.47	88.26	89.29	86.10	74.25	54.76	47.28	60.32
	notMNIST	96.66	67.35	67.35	38.70	96.72	67.66	67.66	43.82	08.95	80.90	80.90	96.04
	omniglot	54.15	94.43	96.66	96.38	57.38	95.57	97.57	97.32	97.21	18.64	05.19	08.02
	CIFAR10-BW	95.01	79.94	79.94	45.68	95.35	82.60	82.60	55.28	16.55	65.71	65.71	99.41
Letters (86% acc)	MNIST	72.58	78.58	81.31	81.34	72.57	77.01	79.94	80.35	85.30	74.11	66.35	67.72
	FMNIST	88.78	71.64	71.64	34.47	87.66	72.07	72.07	39.68	37.67	87.27	87.27	92.98
	KMNIST	87.53	77.76	77.76	66.62	87.40	76.22	76.22	61.94	49.27	78.00	78.00	82.41
	notMNIST	98.48	58.33	58.33	24.81	98.32	54.19	54.19	35.51	06.12	93.57	93.57	98.01
	omniglot	77.47	94.86	97.72	97.88	80.72	95.22	97.83	97.92	86.92	28.63	12.27	10.51
	CIFAR10-BW	98.81	53.28	53.28	07.80	98.88	53.48	53.48	32.55	03.94	97.97	97.97	99.89

5. Results and discussion

This section presents the results obtained from the experiments, along with an analysis that leads to answers to the previously introduced research questions: RQ1 (Section 5.1), RQ2 (Section 5.2) and RQ3 (Section 5.3). The latter includes a short analysis of the limitations that the relevance attribution maps produced by our proposed detector undergo in regards to the complexity of the SNN network structure, as well as the visual features observed by the trained SNN.

5.1. RQ1: Does the proposed SCP-based detector perform competitively when compared to post-hoc OoD detectors from the literature adapted to SNN architectures?

This question will be answered differentiating three types of datasets, each with its own dedicated subsection: Section 5.1.1 for gray-scale datasets, Section 5.1.2 for color datasets and Section 5.1.3 for neuromorphic datasets.

5.1.1. Gray-scale datasets

Scores obtained to address RQ1 are shown in Tables 3 and 4 for the gray-scale datasets and in Table 5 for the color datasets. In all tables, dark gray shaded cells indicate, for each dataset configuration and score, the winning approach among the four OoD detection methods under comparison. Likewise, light gray shaded cells correspond to the second best method in terms of AUROC score (in general the ranking of the methods in the benchmark as per this score is maintained for the rest of scores).

We start by discussing on the results for the gray-scale datasets. Several observations can be done by examining the results corresponding to the FC-SNN model (Table 3). First, the proposed SCP-based OoD detection method achieves in general better detection scores, specially for the MNIST and Letters datasets. Only results obtained for omniglot are worse for the SCP-based approach than for the rest of methods in the benchmark. In the case of the KMNIIST dataset, the SCP-based detector dominates the comparison only when notMNIST and CIFAR10-BW are used as out-distribution datasets. Another interesting fact is that the scenarios where our devised method achieves

the best results, all other logit-based OoD detectors fall apart, scoring significantly worse. Conversely, when any logit-based method performs better than our method, our proposal is never ranked second within the benchmark. This suggests that OoD predictions issued by the SCP-based detector succeed at rectifying those produced by logits-based approaches when they are fused together.

Another observation is that there are some cases where ODIN and Baseline methods perform equally. The reason is that, as explained in Section 4.3, several values of the temperature parameter for each combination of datasets have been tested, reporting in all cases the best score. This could occur for a temperature value equal to 1, which would effectively convert ODIN into the Baseline method. Actually, this occurs in the cases where our method clearly dominates the benchmark, supporting the idea that logits are not suitable for detecting OoD samples in some cases.

We now shift the focus of the discussion on the results for the CNN-SNN architecture shown in Table 4. In this second experiment, our method outperforms the rest of competing OoD detectors in a wider number of cases. Exceptions are Letters and KMNIIST datasets, especially the latter where the SCP-based detector only performs best in 1 out of 6 OoD datasets. Furthermore, performance gaps in those cases where it dominates the rest of the benchmark are small. It is also worth noting that, although our proposal attains superior results in less cases than in the FC-SNN architectures, the differences to the winning OoD technique are not that large in those cases.

The smaller performance gaps observed in this second set of experiments call for a further examination of their statistical significance. To this end, we resort to two different methodologies widely used by the community for the purpose, and apply them to the AUROC scores reported in Tables 3 and 4:

- A critical distance diagram [68], which employs a post-hoc Nemenyi test to compute a critical distance value; if the average rankings of two methods as per the paired score values achieved across different cases are separated by more than the critical distance value, they can be considered to perform significantly different to each other. Otherwise, if their average ranks differ in less than the critical distance, they are declared to be statistically equivalent.

Table 4
OoD detection scores obtained for the gray-scale CNN-SNN architecture.

ID dataset	OoD dataset	AUROC \uparrow				AUPR \uparrow				FPR95 \downarrow			
		SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy
MNIST (97% acc)	FMNIST	92.83	91.64	93.00	91.51	94.53	92.86	93.38	91.92	49.60	54.75	37.87	44.60
	KMNIST	97.23	89.64	90.23	88.33	97.75	90.25	90.51	88.67	13.36	57.13	50.88	58.52
	Letters	92.03	86.04	87.35	86.06	91.19	85.42	86.85	85.75	34.75	63.65	57.05	62.21
	notMNIST	97.10	87.42	88.48	88.04	97.56	86.54	87.60	87.57	17.57	61.43	54.74	58.59
	omniglot	96.91	89.16	90.67	90.06	97.46	88.51	89.78	89.21	15.83	55.73	45.24	48.25
	CIFAR10-BW	94.05	94.83	96.86	95.04	95.84	96.16	97.53	96.19	49.57	45.15	22.29	41.23
FMNIST (77% acc)	MNIST	97.45	49.83	65.76	82.50	97.97	52.79	65.84	80.07	09.34	95.10	88.00	51.39
	KMNIST	97.57	68.51	79.60	88.70	98.09	70.75	80.63	88.06	09.92	87.96	79.97	48.09
	Letters	96.16	53.63	66.31	78.39	96.72	52.99	63.30	74.03	17.64	91.64	85.06	58.87
	notMNIST	95.95	65.24	72.27	73.41	95.52	63.03	68.60	70.38	15.00	87.33	87.46	85.95
	omniglot	97.68	61.10	74.49	86.62	97.67	61.21	72.73	83.80	06.35	91.50	82.96	44.55
	CIFAR10-BW	92.14	83.39	91.23	87.03	92.79	85.41	92.73	89.22	36.89	70.20	61.87	75.06
KMNIST (87% acc)	MNIST	83.96	84.85	86.49	86.43	86.07	87.97	88.91	88.82	56.50	61.79	51.07	53.44
	FMNIST	60.07	74.17	76.94	73.53	63.41	70.51	72.70	70.29	89.34	76.99	66.97	78.09
	Letters	82.26	84.26	87.30	87.01	82.34	86.75	88.98	88.57	55.23	61.34	46.21	48.33
	notMNIST	83.06	84.53	86.99	84.72	84.29	85.52	85.99	83.99	55.66	58.34	42.38	52.86
	omniglot	92.93	84.43	87.78	87.35	93.89	86.33	89.11	88.47	25.60	59.33	43.72	46.16
	CIFAR10-BW	47.35	85.66	92.24	85.36	54.90	89.30	94.10	89.43	99.27	64.13	32.33	69.18
Letters (82% acc)	MNIST	75.90	75.36	75.55	75.49	77.47	79.47	78.71	78.41	81.93	88.74	90.63	91.35
	FMNIST	79.06	76.23	76.23	64.47	84.18	82.24	82.24	69.43	84.80	88.46	88.46	90.42
	KMNIST	86.15	77.05	79.11	74.91	88.61	79.43	81.07	77.74	68.37	82.22	86.67	92.81
	notMNIST	83.35	65.74	65.74	62.31	84.62	66.67	66.67	67.12	66.84	91.88	91.88	96.61
	omniglot	91.57	78.36	84.53	85.18	92.61	80.94	85.64	86.06	43.95	79.14	69.93	68.46
	CIFAR10-BW	67.49	77.56	77.56	49.36	76.24	84.09	84.09	64.09	99.69	93.68	93.68	99.96




Fig. 6. Critical distance diagrams computed over the AUROC results across the cases considered in the experiments for (a) FC- and (b) CNN-SNN architectures. The value of the critical distance is shown on top of every diagram. Horizontal lines connecting the average rank markers of several detectors indicate that their average rankings cannot be claimed to be different to each other with statistical significance.

- A Bayesian posterior analysis of the differences between the paired performance scores achieved by the techniques under comparison. The work in [69] exposed that standard hypothesis testing was not suitable for assessing the significance of performance differences for a number of issues. Alternatively, they proposed to perform a Bayesian sign test of the performance scores of two different techniques, yielding an estimation of the probability that one technique outperforms the other using the scores obtained by each of them over all cases. The probability resulting from this analysis can be carried out via Monte Carlo sampling and depicted in a system of barycentric coordinates, in which three regions are distinguished: one where the first OoD technique outperforms the second, a second one where the second outperforms the first, and a third region reflecting *practical equivalence*. This latter region delimits the probability that techniques can be concluded to perform equivalently from the statistical point of view, and depends on a parameter (*rope*) that establishes the minimum difference between the scores of the OoD techniques for considering their performance as significantly different. The value of *rope* depends on the task being solved: in our case, *rope* is set equal to AUROC = 1%.

Fig. 6 depicts the result of the first statistical significance study over the results obtained for the (a) FC- and (b) CNN-SNN architectures. It is straightforward to note that for both SNN architectures, our method ranks first on average, yet its dominance over the rest of methods cannot be concluded to be statistically significant. Therefore, it is necessary to perform a closer study between our SCP-based detector and the second best one, namely, ODIN.

This closer examination is done by performing a Bayesian analysis of the differences following the methodology described above using the AUROC scores of ODIN and the proposed SCP-based detector. Figs. 7a and 7b show the results of this second significance study. When it comes to the scores obtained for the FC-SNN architecture (Fig. 7a), almost all Monte Carlo samples drawn from the adjusted Bayesian posterior probability lie in between the regions where ODIN or the SCP-based detector achieve better results. This means that none of the methods perform equivalently, and one always outperforms the other with statistical significance. This goes in line with the observation made in the discussions on Table 3, where it was concluded that either a logit-based approach or the proposed SCP-based detector performs best in the majority of cases, being very few the scenarios where both perform similarly. The Bayesian analysis reported in Fig. 7a informs that the probability that our method performs better than ODIN is approximately 73%, whereas ODIN is expected to perform better for the remaining percentage.

We follow this second significance analysis by exploring the output of the Bayesian analysis over the results obtained for the CNN-SNN architecture. Fig. 7b clearly reveals that either our method performs better (approximately 94%) or ODIN does (6%), whereas the probability that both perform equivalently (*rope* region) is almost negligible.

5.1.2. Color datasets

Regarding the color datasets, the results are more evenly matched as exposed by the scores reported in Table 5. However, conclusions obtained are similar to the ones drawn for the experiments with gray-scale datasets.

Table 5
OoD detection scores obtained for the RGB CNN-SNN architecture.

ID dataset	OoD dataset	AUROC ↑				AUPR ↑				FPR95 ↓			
		SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy
CIFAR10 (65% acc)	Food101	71.65	62.69	68.10	67.45	70.30	64.06	70.12	69.80	81.56	91.49	89.99	92.31
	LSUN (crop)	87.63	65.91	74.09	74.67	86.97	70.84	77.93	78.40	49.06	93.14	90.15	92.42
	LSUN (resize)	69.54	65.65	72.80	74.46	68.95	67.01	74.21	74.97	84.18	89.11	86.81	84.59
	Flowers102	75.77	62.79	68.71	67.75	75.04	64.71	70.93	70.32	75.59	92.47	90.54	91.84
	Caltech101	67.02	61.79	65.88	66.47	64.47	63.75	67.90	68.58	83.81	92.52	91.53	91.79
	GTSRB	78.76	67.03	71.81	69.12	78.10	70.41	75.42	73.61	69.49	92.11	91.38	91.71
	SVHN	62.70	66.73	70.10	69.72	61.21	69.28	72.27	72.06	88.51	90.50	89.41	90.17
	TinyImagenet	54.03	70.89	70.89	69.61	55.63	75.39	75.39	74.70	95.34	90.44	90.44	93.02
	EuroSAT	65.88	66.49	70.00	69.56	64.91	68.73	71.43	70.70	85.69	90.04	88.96	89.60
	DTD	61.43	76.86	81.76	73.19	63.91	80.22	85.12	83.92	90.60	84.28	83.95	80.71
SVHN (75% acc)	Food101	61.73	61.95	63.97	65.64	59.24	61.54	64.94	67.58	84.89	91.49	92.71	92.55
	LSUN (crop)	84.74	76.35	74.84	80.20	85.82	78.02	81.40	83.76	50.59	90.97	94.20	80.90
	LSUN (resize)	73.20	77.39	81.73	82.80	76.87	90.41	85.49	86.29	73.95	86.09	86.75	77.49
	Flowers102	80.47	74.29	77.92	81.32	81.66	79.00	82.68	84.66	60.79	90.28	93.20	78.76
	Caltech101	80.94	73.85	77.24	78.39	82.75	78.50	81.86	82.42	60.35	89.45	91.97	84.31
	GTSRB	81.82	73.74	77.02	79.88	82.23	78.34	81.85	83.45	56.24	93.60	90.09	82.10
	CIFAR10	83.11	78.02	81.01	78.90	83.80	77.85	81.10	81.79	55.59	86.42	88.64	81.76
	TinyImagenet	80.23	73.73	77.80	80.90	81.23	78.67	82.71	84.46	60.81	90.93	93.08	81.21
	EuroSAT	79.06	74.02	77.73	80.47	79.38	78.53	82.25	83.87	60.84	89.79	92.32	80.48
	DTD	81.16	74.11	77.69	81.01	92.85	78.78	82.39	84.42	56.75	81.61	81.61	78.89

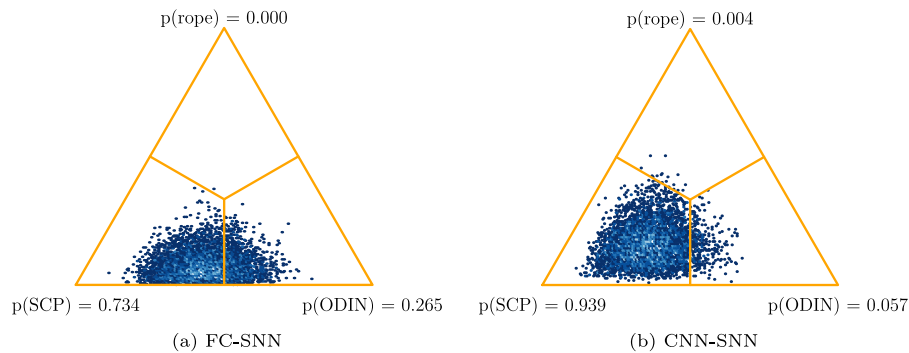


Fig. 7. Posterior probability resulting from the Bayesian signed tests computed over the paired AUROC scores attained by ODIN and the proposed SCP-based detector with (a) the FC-SNN architecture; (b) the CNN-SNN architecture.




Fig. 8. Critical distance diagram computed over the AUROC results across the cases considered in the experiments for color datasets.

In general, the detection performance of all methods is worse than in the previous set of experiments, due to the fact that the accuracy of the model is significantly lower (approximately 65% and 70% for CIFAR10 and SVHN, respectively). In these two in-distribution datasets, SCP outperforms the other detection methods over 5 out of the 10 out-distribution datasets. This can be further observed in Figs. 8 and 9. In the critical distance diagram, where obtained average rank is taken into account, the winning method is (Free) Energy, but without statistical significance as the distance of its average rank to that of ODIN and the proposed SCP is smaller than the critical distance. Therefore, further exploration is carried out by means of a Bayesian analysis, concluding that our method works significantly better than ODIN (approximately 82% of the times), while performing on par with Free Energy, as both perform better than the other with equal likelihood (45%).

5.1.3. Neuromorphic datasets

Finally, we analyze the results obtained for the neuromorphic datasets, which are summarized in Table 6. It can be observed that the SCP-based detector clearly outperforms its competitors in all the ID vs OoD combinations, except for the DVSGesture vs MNIST, where all datasets attain an optimal detection performance. In fact, our SCP based detector obtains near-perfect scores (more than 99% AUROC) when MNIST is the in-distribution dataset.

Results are in line with expectations, as SNNs are by design well suited for processing and modeling temporal data. Therefore, samples of these neuromorphic datasets can be characterized more clearly in the activations space than in the logit space, giving rise to a better latent embedding for the spike count characterization performed by SCP.

5.2. RQ2: How does the SCP-based detector perform when fused with other adapted OoD detection methods?

To inform this research question with evidence, different ensembles have been built by fusing the OoD detection methods considered for the benchmark. At this point it is important to note that the fusion of several techniques is only worth when (i) they yield similar levels of performance for the datasets being evaluated; and (ii) when their decisions for a given input are *diverse*. Otherwise, poorly performing techniques can introduce noise to the final fused decision and ultimately, degrade the overall performance of the ensemble. Therefore, ensembles will be assessed for the color datasets, where the OoD methods to be fused perform most closely to each other.

Table 6
OoD detection scores obtained for the RGB CNN-SNN architecture in neuromorphic datasets.

ID dataset	OoD dataset	AUROC \uparrow				AUPR \uparrow				FPR95 \downarrow			
		SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy
NMNIST (97% acc)	DVSGesture	99.76	83.29	83.29	81.96	99.80	87.39	87.39	85.18	00.00	84.77	84.77	80.50
	CIFAR10DVS	99.89	88.86	92.28	92.17	99.90	90.16	92.9	92.96	00.00	66.15	42.51	43.01
	DVSLip	99.12	92.05	92.05	89.81	99.42	93.44	93.44	90.56	01.12	54.13	54.13	50.07
DVSGesture (70% acc)	NMNIST	100.0	100.0	100.0	90.72	100.0	100.0	100.0	95.25	00.00	00.00	00.00	43.01
	CIFAR10DVS	85.74	53.89	66.56	57.85	87.35	56.36	65.95	58.63	24.62	82.20	71.21	85.98
	DVSLip	99.20	81.27	95.02	83.41	98.32	77.84	93.40	84.08	00.76	36.36	11.74	58.33




Fig. 9. Posterior probability resulting from the Bayesian signed tests computed over the paired AUROC scores attained by ODIN, Energy and the proposed SCP-based detector.




Fig. 10. Fusion of OoD detection methods and voting criterion for their predictions.




Fig. 11. Critical distance diagram of ensembles, ODIN and Free Energy evaluated over color datasets.

Fig. 10 illustrates how the fusion of two OoD detection methods A and B works: for a given query instance X_{test} , the predictions issued by both fused methods are combined together under an OR criterion, so that the outcome of the ensemble for the query instance is OoD (out of distribution) only if both methods predict the instance as such. Otherwise, the sample is predicted to belong to the in-distribution dataset. By building ensembles with different methods, we aim to elucidate whether the two methods perform diversely yet coherently enough for their combination to yield better detection scores w.r.t the performance of the two methods in isolation.

When addressing RQ2, several ensembles are created, namely, *E-ODIN-Energy*, *E-ODIN-SCP* and *E-Energy-SCP*, and evaluated over color datasets, where differences reported in response to RQ1 are narrower. If the fusion of models produce similar decisions for the test samples in the different experiments, the ensemble will not improve the scores achieved by its fused methods. On the other hand, if decisions by the combined methods differ from each other, the fusion may yield differences in detection performance. Table A.1 included in Appendix A

comprehensively collects the detection performance figures for these ensembles. For the sake of clarity, we herein discuss on the CD diagram (Fig. 11) and the barycentric plot resulting from the Bayesian analysis of performance differences (Fig. 12).

We begin by inspecting the critical distance diagram, which clearly depicts that both Free Energy and ODIN (logits-based approaches) perform outstandingly better when fused with the SCP-based detector. Contrarily, when Free Energy and ODIN are fused together to yield the *E-ODIN-Energy* ensemble, the fusion of predictions seem not to lead to significant performance gains. Actually, all three methods' and the *ODIN-Energy* ensemble's average ranks are separated by less than the critical distances, meaning that their average ranks are not statistically different from each other. As in RQ1, this lack of statistical significance calls for a deeper analysis towards solidly answering RQ2.

When applying a Bayesian analysis over the obtained results, we take into account not only the difference in ranks, but also the difference in raw performance, using AUROC as the metric like in the previous cases. In order to make a fair comparison, we compare the *E-ODIN-SCP* and *E-ODIN-Energy* ensembles to the Free Energy method, to verify whether the fusion of SCP with ODIN is more beneficial than combining ODIN with Free Energy. Such comparisons can be observed in Figs. 12a and 12b, which clearly show how the fusion with our SCP-based technique outperforms Free Energy almost all the times. Specifically, Fig. 12a demonstrates that the *E-ODIN-SCP* ensemble, despite using ODIN (a lower performing method when compared to Free Energy), turns out to achieve better results with very high likelihood (99.6%). Conversely, the *E-ODIN-Energy* ensemble is not even capable of outperforming Free Energy, one of the techniques that composes it. Together with the detailed report of scores in Table A.1, it can be concluded that the predictions issued by our SCP-based method, when fused with those made by other OoD detectors, allow detecting OoD instances more effectively than any other approach in the benchmark, either individual methods or any other ensemble not considering the SCP-based detector at its core.

5.3. RQ3: Does the local relevance attribution method yield informative explanations about the features that make a given sample be detected as ood? Which are the limitations of this explainability technique?

At this point we recall that the local explanation technique for the proposed SCP-based detector aims to generate a relevance attribution




Fig. 12. Posterior probability resulting from the Bayesian signed tests computed over the paired AUROC scores attained by the *E-ODIN-SCP* ensemble, the *E-ODIN-Energy* ensemble and the Free Energy method.

Table 7
Scores for the SCP-based detector in the experiments devised to address RQ3.

ID dataset	OoD dataset	Model	AUROC \uparrow	AUPR \uparrow	FPR80 \downarrow
MNIST	MNIST-Square	FC-SNN	96.04	96.70	02.92
		CNN-SNN	69.73	69.06	55.14
	MNIST-C zigzag	FC-SNN	83.73	83.03	27.05
		CNN-SNN	75.06	73.64	45.11

vector $\mathbf{h}(\mathbf{x})$ for a query sample \mathbf{x} (i.e. a heatmap when dealing with image classification). The value of its components $\{h_d(\mathbf{x})\}_{d=1}^D$ indicates the relative relevance of regions of the input space that have influenced most the detector when detecting that the query sample is OoD. The heatmap must also deliver information about an *OoD score* of the sample, such that the intensity of the heatmap should be higher as the OoD score of the sample increases, and vice-versa.

Based on the above intuition, the qualitative evaluation of the attribution method made to reply RQ3 proceeds as follows: first, attribution generated for several test images with known artifacts are shown and analyzed to verify that the produced heatmaps spot those artifacts. Next, limitations of the attribution method noted during the experimentation are exposed, using examples and results that clearly illustrate these constraints.

We concentrate the discussion on FC-SNN and CNN-SNN architectures trained to classify MNIST images (in-distribution dataset), using query instances from MNIST-Square and MNIST-C as out-of-distribution datasets. The detection performance of the SCP-based OoD detector in this setup is reported in Table 7, indicating values of the three detection scores used so far. We note that in this case the value of FPR80 is reported, i.e., the FPR value obtained when the target TPR is set to 80%. In other words, the detection threshold is set such that 80% of the positive (in-distribution) samples are correctly detected as such. It is important to note that all the other three logit-based detectors considered in the benchmark obtain worst detection scores in this setup.

After training both SNN architectures over the MNIST dataset, several in-distribution and out-of-distribution samples are tested and processed through the SCP-based detector and the local relevance attribution technique. Fig. 13 depicts several of such query images. The text at the bottom of each image represents the predicted class for the image, whereas the number above the image corresponds to the difference between the OoD score for the given sample – given by the value returned by $f_{OoD}(\mathbf{x}, \hat{y})$ as per Expression (4) – and the threshold $\lambda^{\hat{y}}$ for a TPR of 80%. Hence, a positive value above every image means that the image is detected as OoD. Conversely, the more negative this value is, the more in-distribution the sample can be thought to be, as it is *farther* from being detected as OoD based on the criterion in Expression (5).

The discussion starts with the images shown for the FC-SNN architecture (Fig. 13a). The first example, a digit corresponding to class 5, the rightmost tip of the digit is highlighted as relevant even if clearly belonging to the in-distribution. Indeed, despite its largely negative

score (therefore being confidently detected as ID), usually the 5 digits in the MNIST dataset do not reach the edges of the image. In this case, the 5 sample depicted in the figure has an untypically long tip. In the square and zigzag case this observation is further buttressed: not only the same digit with square and zigzag artifacts preserve a highlighted right tip, but relevance is also attributed by our approach to the parts of the image affected by the artifacts. The 5 digit with the zigzag pattern is not be detected as OoD, but its OoD score is close to 0, hence the image is more likely to be an OoD than the normal case. Consequently, the produced heatmap is less intense. The 8 digit comprising square and zigzag artifacts would be detected as OoD, an both artifacts are partly remarked in the heatmaps. Some parts of the clean 8 digit image are also marked in the heatmap: this can be explained because the closest cluster for the samples with the corruption changes w.r.t. the clean digit from where they are produced, which may be composed by digits with different compositional characteristics. In the last example, the 4 digit, the zigzag case is predicted to be a digit 9 rather than a digit 4, and thus highlights several regions where the shape of an archetypical 9 could be.

Digits shown for the CNN-SNN architecture exhibit in general noisier and smoother heatmaps, as relevance attribution vectors at the input of the architecture are scaled up via interpolation to match the dimensions of the original input image in this case, from 11×11 to 28×28 pixels. In all examples included in Fig. 13b, the artifacts induced in the clean 1, 9 and 7 digits are highlighted in the heatmap. In the case of the digit 7, the noisier heatmaps are due to the fact that the predicted classes are 9 (square) and 2 (zigzag) rather than a 7. Hence, the parts of the image spotted by the heatmaps include both the artifacts and the parts of the original digit that make the spike count patterns deviate from those of the closest centroid of its wrongly predicted class. In summary, we conclude that the proposed relevance attribution technique can generate heatmaps that provide meaningful information about the features of the input that make the SCP detector identify it as an out-of-distribution sample. However, during our experiments we noted several limitations of the approach that spur future research lines later described in the concluding part of the manuscript. Such limitations are:

- *The depth of the network degrades the quality of attribution heatmaps:* one effect observed when applying the proposed relevance attribution technique is that heatmaps become noisier as the depth of the fully connected part of the SNN increases, i.e., the number of fully connected layers L . This effect can be observed in Fig. 14,

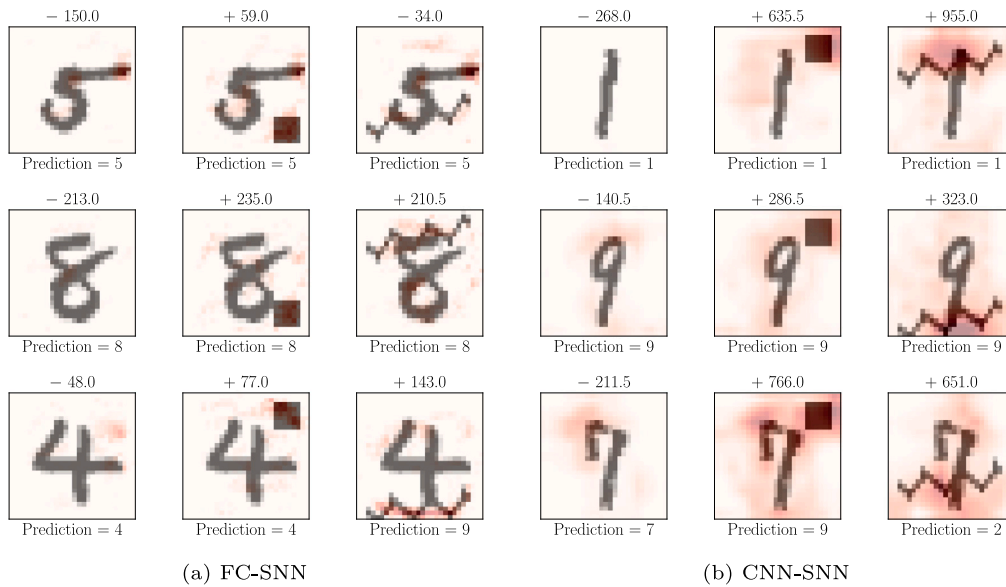


Fig. 13. Heatmaps showing the relevance attributed to different digits (clean and corrupted versions with square and zigzag artifacts) for (a) a FC-SNN architecture; (b) a CNN-SNN architecture. Each heatmap is superimposed on its input, and for every image its predicted class is reported, together with the difference between its OoD score and the class-conditional threshold λ^c .

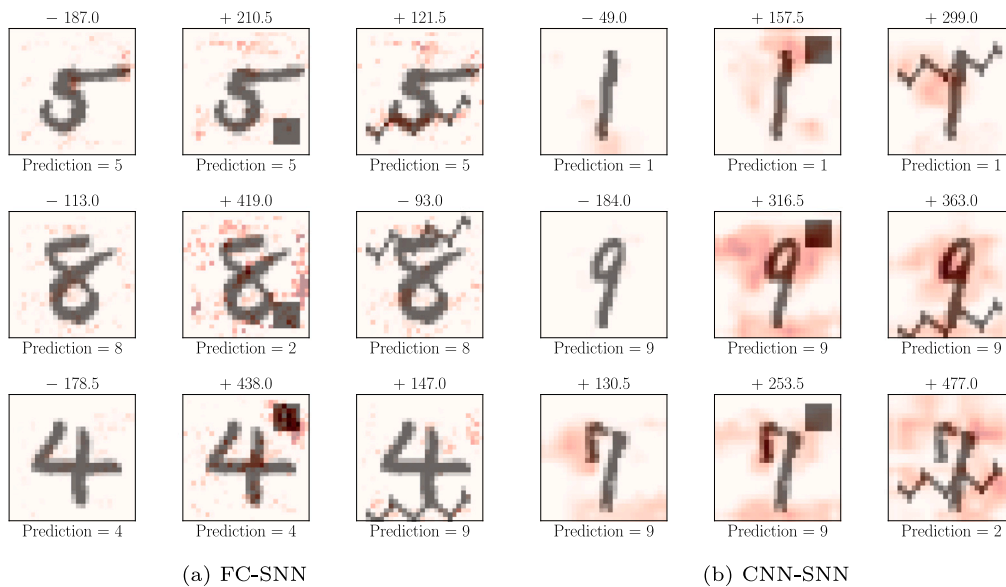


Fig. 14. Heatmaps showing the relevance attributed to different digits (clean and corrupted versions with square and zigzag artifacts) for (a) a FC-SNN architecture; (b) a CNN-SNN architecture. In both cases the number of fully connected SNN layers is $L = 2$. The layout and interpretation of the plots are identical to those in Fig. 13.

where the same inputs as in Fig. 13 are used over FC-SNN and CNN-SNN architectures with $L = 2$ fully-connected layers, instead of the single-layered architecture considered in previous experiments.

As noted in these plots, in general the attribution technique elicits less sharply defined heatmaps, evincing its lower capability to emphasize important features in the input image towards the OoD detection. Nonetheless, for some samples their predicted class has changed. Specifically, the digit 8 with the square artifact in the FC-SNN model, and the digit 7 in the CNN-SNN architecture. Moreover, the lower definition and quality of the produced attributions is not reflected on the detection scores, as can be verified in Tables B.1 and B.2 and in Fig. B.1 of Appendix B. While in the case of the CNN-SNN architecture a small performance loss is noted, when it comes to the FC-SNN architecture the

developed SCP-based detector yields even better detection scores when compared to ODIN.

- *The attribution depends on features extracted from the input by the network:* Our detector stringently depends on the features extracted from the input image and that are input to the fully-connected part of the SNN. That is, when it is preceded by a feature extractor like a convolutional layer, our detection capabilities depend on the ability of this layer to learn feature maps that capture characteristics that push the SCP-detector to declare that the image is out-of-distribution. Unfortunately, satisfying this necessary condition is not sufficient to guarantee that OoD samples will be detected as such. A false negative could lead to a heatmap that does not highlight characteristics that expose the OoD nature of the misdected input sample. On the one hand, as feature maps go through a fully connected network before reaching the layer where the SCP-based detection

is performed (the layer prior to softmax), the effect on the latter could be compensated for by the learned weights, as neurons are densely connected and every neuron affects all the neurons of the subsequent layers. On the other hand, due to the way the OoD score is computed, a high intra-cluster variance could eventually mask an OoD sample. As described in Algorithm 1, our method assigns an OoD score computed as the L_1 distance to the closest centroid among the spike count clusters of the predicted class. Hence, if the spike count pattern of the input changes, the closest cluster to the input may vary accordingly, specially if the intra-cluster variance is high. This is not a bad effect overall, as it allows in-distribution samples to be more accurately represented by the centroids. But what could also occur is that the reconstruction of one of the centroids (mentioned in Section 3.3) of the predicted class coincidentally matches the reconstruction of the input being tested, masking the OoD characteristics that should be highlighted as relevant.

Notwithstanding the above limitations, results reported in this section support that (i) the proposed SCP-based OoD detector performs very competitively with respect to other model-agnostic OoD detection approaches from the literature, with significance as per a Bayesian analysis of the performance differences; and (ii) that the relevance attribution technique effectively highlights artificially induced modifications to in-distribution data that a priori should contribute to their identification as OoD by the SCP-based detector.

6. Conclusions and future work

This work has elaborated on the out-of-distribution detection problem when dealing with spiking neural network architectures. Focusing on networks with rate-based encoding schemes and leaky integrate-and-fire neurons, the main idea of the proposed OoD detector is to leverage the spike count patterns produced at the last intermediate layer of the neural model to create archetypes or representatives of those patterns for every class in the dataset at hand. By comparing these archetypes with spike count patterns stimulated by new query samples, the proposed detector can effectively discriminate whether they are in- or out-distributions samples. Comparisons are made by using a distance metric, whereas archetypes are represented by the centroids of a cluster space computed over the spike count patterns of training samples by a clustering approach. We further inform the output of the OoD detector with explanations issued by a local relevance attribution method, providing a measure of the features of any given input sample that drive most the output of the developed detector towards declaring the input as out-of-distribution.

Extensive experiments done for image classification tasks over fully-connected and convolutional-based SNN architectures, for gray-scale, color and event-based datasets, have verified that the proposed detector achieves a competitive detection performance over other model-agnostic post-hoc OoD detection techniques, which have been tailored to work with spiking-based neural models. Results are specially impressive in the neuromorphic datasets, where our technique consistently outperforms competitors. In addition, experiments with ensembles of OoD detection techniques have showcased that the fusion of the proposed SCP-based detector with a logits-based method attains unrivaled levels of performance when compared to other combinations of techniques. Furthermore, explanations issued by the proposed local relevance attribution method excel at identifying artificially inserted perturbations on in-distribution samples, which are highlighted as the parts of the image that are influential for their detection as OoD. Limitations of this technique have been exposed and discussed, showing (i) the degradation of the quality of explanations with the increasing depth of the fully connected part of the network; (ii) the dependence of the attributed relevance on the input features captured by the input of the SNN model; and (iii) the intra-cluster variance of clusters computed from the spike counts of the input data.

Several research directions are envisioned for the near future arising from the promising results presented in this work. To begin with, we plan to extend the applicability of the SCP detector and the principles on which its design relies to other types of encoding methods suitable for spiking neural architectures, especially those affecting most the spike count strategy that lies at the core of the SCP based detector, e.g. the so-called time-to-first-spike encoding (TTFS [70]). Efforts will be also invested towards mitigating the impact of the aforementioned limitations on the produced heatmaps of the relevance attribution technique. Other forms of explanatory information about the detection of OoD samples will be explored, such as the use of generative spiking models for counterfactual studies of the audited SNN model. In addition, we plan to take a step beyond the fusion of OoD methods to combine them with the output of uncertainty estimation techniques, which can also give a hint about the unknown nature of a given query to the model. Finally, OoD detection in temporal or spatio-temporal streaming data, such as brain EEG and fMRI data [71], seismic data [29] or financial data [72] is another direction that we plan to follow in future studies.

CRediT authorship contribution statement

Aitor Martinez-Seras: Conceptualization, Investigation, Data curation, Methodology, Writing – original draft, Writing – review & editing, Software. **Javier Del Ser:** Conceptualization, Investigation, Validation, Visualization, Writing – original draft, Writing – review & editing, Supervision, Funding acquisition. **Jesus L. Lobo:** Validation, Writing – review & editing. **Pablo Garcia-Bringas:** Writing – review & editing, Supervision, Funding acquisition. **Nikola Kasabov:** Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

A github repository is linked in the experimental setup section of the paper.

Acknowledgments

A. Martinez Seras receives funding support from the Basque Government through its BIKAINTEK PhD support program. J. Del Ser acknowledges funding support from the same institution through the Consolidated Research Group MATHMODE (IT1456-22) and the ELKARTEK program (EGIA, grant no. KK-2022/00119, and BEREZ-IA, ref. KK-2023/00012). P. García Bringas also thanks the funding support from the Basque Government through the Consolidated Research Group D4K - Deusto for Knowledge (IT1528-22) and the ELKARTEK funding grants REMEDY (ref. KK-2021/00091) and SIIRSE (ref. KK-2022/00029).

Appendix A. Results of ensembles in color datasets

See Table A.1.

Appendix B. OoD detection results for SNN with increased depth

See Tables B.1, B.2 and Fig. B.1.

Table A.1

OoD detection scores obtained by the individual methods and the ensembles over the color datasets using a convolutional SNN architecture.

ID dataset	OoD dataset	AUROC ↑						AUPR ↑						FPR95 ↓					
		SCP (proposed)	ODIN	Energy	Ensemble ODIN-SCP	Ensemble ODIN-Energy	Ensemble Energy-SCP	SCP (proposed)	ODIN	Energy	Ensemble ODIN-SCP	Ensemble ODIN-Energy	Ensemble Energy-SCP	SCP (proposed)	ODIN	Energy	Ensemble ODIN-SCP	Ensemble ODIN-Energy	Ensemble Energy-SCP
CIFAR10 (65% acc)	Food101	71.65	68.10	67.45	74.85	67.75	74.56	70.30	70.12	69.80	74.74	69.96	74.64	81.56	89.99	92.31	98.13	95.24	98.55
	LSUN (crop)	87.63	74.09	74.67	86.34	74.96	86.70	86.97	77.93	78.40	87.30	78.64	87.52	49.06	90.15	92.42	95.11	96.20	96.06
	LSUN (resize)	69.54	72.80	74.46	76.43	73.58	76.86	68.95	74.21	74.97	76.17	74.58	76.23	84.18	86.81	84.59	97.90	91.86	97.56
	Flowers102	75.77	68.71	67.75	77.36	68.04	76.82	75.04	70.93	70.32	77.73	70.58	77.42	75.59	90.54	91.84	97.76	95.72	98.26
	Caltech101	67.02	65.88	66.47	70.40	65.96	70.81	64.47	67.90	68.58	69.66	68.11	70.02	83.81	91.53	91.79	98.61	95.62	98.64
	GTSRB	78.76	71.81	69.12	81.58	69.90	80.07	78.10	75.42	73.61	82.31	74.25	81.49	69.49	91.38	91.71	97.50	95.87	97.47
	SVHN	54.03	70.89	69.61	67.26	71.91	67.06	55.63	75.39	74.70	68.13	76.80	68.18	95.34	90.44	93.02	99.65	97.54	99.69
	TinyImagenet	65.88	70.00	69.56	72.59	69.76	72.26	64.91	71.43	70.70	72.20	71.05	71.73	85.69	88.96	89.60	98.36	94.06	98.52
	EuroSAT	61.43	81.76	81.01	76.89	81.36	76.44	63.91	85.12	83.92	77.94	84.56	77.70	90.60	83.95	80.71	98.38	91.06	98.08
	DTD	61.73	63.97	65.64	65.38	64.46	66.73	59.24	64.94	67.58	64.68	65.81	66.30	84.89	92.71	92.55	98.62	96.60	98.67
SVHN (75% acc)	Food101	84.74	76.35	80.20	88.08	77.35	89.64	85.82	81.40	83.76	89.79	82.10	90.78	50.59	94.20	80.90	97.11	94.69	90.72
	LSUN (crop)	73.20	81.73	82.80	85.96	81.63	86.23	76.87	85.49	86.29	88.03	85.52	88.24	73.95	86.75	77.49	96.53	89.26	94.12
	LSUN (resize)	80.47	77.92	81.32	87.31	78.79	88.43	81.66	82.68	84.66	88.65	83.28	89.29	60.79	93.20	78.76	97.36	94.00	91.68
	Flowers102	80.94	77.24	78.39	87.01	77.62	87.29	82.75	81.86	82.42	89.01	82.12	89.12	60.35	91.97	84.31	96.55	93.48	93.67
	Caltech101	81.82	77.02	79.88	86.77	77.82	87.77	82.23	81.85	83.45	87.90	82.38	88.46	56.24	93.60	82.10	97.25	94.48	91.99
	GTSRB	83.11	78.02	78.90	87.47	78.39	87.95	83.80	81.10	81.79	88.04	81.55	88.43	55.59	88.64	81.76	94.80	90.63	91.93
	CIFAR10	80.23	77.80	80.90	86.98	78.67	88.01	81.23	82.71	84.46	88.38	83.31	88.92	60.81	93.08	81.21	97.25	94.03	92.46
	TinyImagenet	81.16	77.69	81.01	86.92	78.55	88.10	81.61	82.39	84.42	88.09	83.01	88.79	56.75	92.85	78.89	97.06	93.68	90.69
	EuroSAT	42.77	90.31	88.85	60.47	89.91	60.18	47.66	91.80	90.96	60.29	91.58	60.04	95.18	59.33	63.11	97.92	68.55	98.00
	DTD	58.09	83.81	84.08	72.74	84.01	72.65	57.29	86.87	87.15	69.95	87.07	69.82	83.03	78.72	73.67	95.80	82.66	94.95

Table B.1

OoD detection scores corresponding to the FC-SNN architecture with $L = 2$ fully connected layers.

ID dataset	OoD dataset	AUROC ↑				AUPR ↑				FPR95 ↓			
		SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy
MNIST (97% acc)	FMNIST	85.76	85.57	87.16	86.40	86.79	83.79	86.59	87.27	60.73	58.38	59.08	67.25
	KMNIST	91.43	91.41	92.10	91.19	91.74	91.79	92.17	91.33	39.63	48.81	42.97	47.24
	Letters	87.96	87.26	87.99	86.89	87.80	86.20	86.97	86.37	49.71	58.87	55.34	62.66
	notMNIST	97.41	82.30	82.30	77.11	97.76	80.49	80.49	74.44	10.71	68.61	68.61	79.02
	omniglot	81.83	96.01	97.97	98.34	82.24	96.68	98.23	98.58	72.01	27.77	11.60	08.72
	CIFAR10-BW	92.26	85.46	86.20	84.69	93.21	86.86	87.34	86.13	39.57	70.46	69.40	75.46
FMNIST (83% acc)	MNIST	94.60	83.26	93.70	95.69	93.45	85.73	94.26	95.85	20.57	66.17	33.27	21.27
	KMNIST	89.81	77.10	85.24	86.45	89.41	79.26	85.58	86.66	45.26	77.71	64.04	61.74
	Letters	94.06	81.00	90.97	92.66	94.14	83.75	91.86	93.23	28.31	71.37	46.87	38.29
	notMNIST	94.29	67.69	77.01	80.85	95.16	71.63	78.21	81.14	34.62	85.43	78.45	74.44
	omniglot	81.58	88.06	95.67	97.28	77.71	88.51	95.56	97.28	59.30	50.79	20.41	12.00
	CIFAR10-BW	79.90	52.31	68.27	75.53	83.78	64.79	74.69	80.03	85.64	96.01	92.53	90.52
KMNIST (89% acc)	MNIST	77.28	83.84	85.35	84.89	78.85	86.20	87.26	86.96	69.90	53.45	44.99	47.75
	FMNIST	78.36	84.41	84.41	77.08	80.04	87.68	87.68	81.86	68.49	55.48	55.48	71.94
	Letters	82.23	84.99	86.54	85.87	83.04	87.70	88.73	88.24	59.82	53.03	43.49	46.78
	notMNIST	95.68	79.61	79.61	75.00	96.21	82.14	82.14	78.10	12.30	63.42	63.42	75.35
	omniglot	72.45	92.49	95.45	96.29	74.39	94.01	96.43	97.16	78.65	25.38	09.77	06.99
	CIFAR10-BW	91.66	89.67	90.74	87.95	92.77	91.92	92.70	91.14	30.70	36.63	27.44	41.20
Letters (89% acc)	MNIST	74.95	80.32	82.40	82.49	73.98	78.76	79.47	79.39	78.48	71.56	61.60	62.14
	FMNIST	87.85	81.80	81.80	66.07	88.17	83.07	83.07	70.52	47.97	76.42	76.42	88.31
	KMNIST	88.80	82.74	84.38	81.75	89.23	83.35	84.28	81.92	48.59	71.89	67.89	73.30
	notMNIST	98.12	74.57	74.57	66.70	98.14	72.27	72.27	64.35	07.90	80.19	80.19	82.95
	omniglot	84.65	93.75	97.66	97.84	86.91	94.23	97.88	98.05	73.25	33.94	12.99	11.47
	CIFAR10-BW	98.21	71.32	71.32	50.54	98.41	75.89	75.89	60.61	07.46	91.11	91.11	98.78

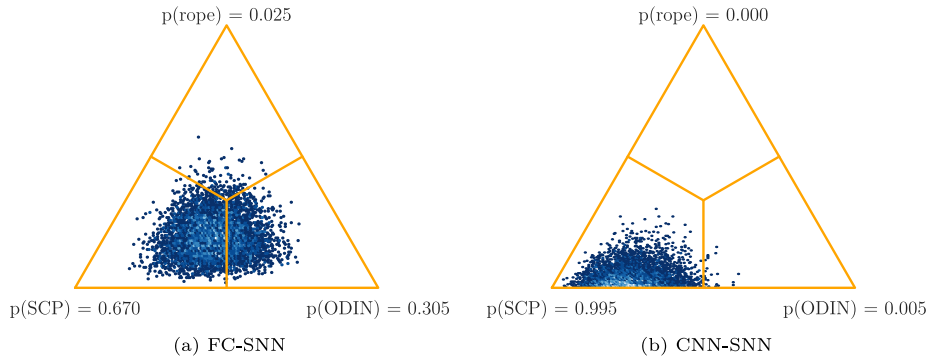


Fig. B.1. Posterior probability resulting from the Bayesian signed tests computed over the paired AUROC scores attained by ODIN and the proposed SCP-based detector with (a) the FC-SNN architecture; (b) the CNN-SNN architecture. In both cases $L = 2$.

Table B.2
OoD detection scores corresponding to the CNN-SNN architecture with $L = 2$ fully connected layers.

ID dataset	OoD dataset	AUROC \uparrow				AUPR \uparrow				FPR95 \downarrow			
		SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy	SCP (proposed)	Baseline	ODIN	Energy
MNIST (98% acc)	FMNIST	95.47	88.32	89.85	88.52	95.96	88.22	89.40	88.22	21.28	53.85	44.14	49.56
	KMNIST	95.55	90.83	91.70	91.28	96.02	90.62	91.13	90.65	20.42	46.69	39.71	40.91
	Letters	90.18	83.60	84.19	83.87	88.26	78.88	78.46	78.20	35.61	54.41	50.13	51.22
	notMNIST	96.27	85.94	86.80	86.79	96.70	83.07	84.11	84.11	16.44	53.70	50.28	51.14
	omniglot	95.59	91.67	93.18	93.36	96.10	91.81	93.00	93.16	20.81	42.94	33.48	32.51
	CIFAR10-BW	97.09	94.66	96.03	94.96	97.72	95.20	96.20	95.19	13.20	30.87	19.62	26.71
FMNIST (82% acc)	MNIST	94.33	78.04	90.45	92.96	94.37	80.92	91.15	93.30	24.22	72.16	41.28	31.64
	KMNIST	92.76	78.67	88.27	90.85	93.29	81.05	88.55	90.68	34.72	70.88	47.62	37.46
	Letters	91.38	76.29	88.48	90.58	91.31	79.63	89.54	91.25	38.28	75.29	48.55	41.47
	notMNIST	93.68	77.25	86.14	86.14	93.44	77.25	83.64	83.31	26.95	74.95	52.48	52.68
	omniglot	94.68	79.66	89.85	92.47	94.96	82.78	90.65	92.90	24.31	73.41	45.13	35.59
	CIFAR10-BW	85.68	82.16	88.50	89.58	85.63	84.33	89.08	90.02	58.43	66.38	48.84	44.91
KMNIST (92% acc)	MNIST	89.19	88.92	90.57	90.49	90.80	90.51	91.80	91.69	28.61	31.88	23.16	23.79
	FMNIST	83.02	88.06	91.24	91.86	86.92	90.95	93.24	93.66	54.23	40.40	24.22	22.84
	Letters	92.07	88.72	90.81	90.95	93.61	90.61	92.18	92.28	18.26	34.76	23.09	22.90
	notMNIST	90.85	88.20	89.45	89.22	92.41	90.35	91.18	90.91	23.56	36.08	28.65	30.06
	omniglot	91.15	89.72	92.38	93.07	93.07	91.48	93.66	94.24	22.17	30.11	17.65	15.68
	CIFAR10-BW	81.63	93.75	96.08	96.79	85.74	95.12	97.12	97.78	54.67	12.91	04.54	03.16
Letters (90% acc)	MNIST	76.04	81.15	84.48	84.38	73.44	81.31	83.94	84.05	74.44	73.56	62.23	64.79
	FMNIST	87.12	85.19	85.19	74.39	88.25	87.25	87.25	75.27	61.90	70.03	70.03	74.37
	KMNIST	91.26	85.31	87.47	84.11	91.90	85.88	87.09	83.60	44.83	66.51	55.99	64.45
	notMNIST	87.38	76.31	76.44	72.06	87.12	76.46	75.96	71.64	51.15	80.72	77.84	81.81
	omniglot	93.32	87.76	92.07	91.95	93.99	88.02	91.44	91.17	36.86	58.40	37.39	36.73
	CIFAR10-BW	96.88	88.67	88.67	73.60	97.56	91.55	91.55	81.54	17.69	73.17	73.17	97.13

References

[1] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, et al., Highly accurate protein structure prediction with AlphaFold, *Nature* 596 (7873) (2021) 583–589.

[2] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *NeurIPS* 33 (2020) 1877–1901.

[3] A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, et al., Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, *Inf. Fusion* 58 (2020) 82–115.

[4] A. Adadi, M. Berrada, Peeking inside the black-box: A survey on explainable artificial intelligence (XAI), *IEEE Access* 6 (2018) 52138–52160.

[5] E. Tjoa, C. Guan, A survey on explainable artificial intelligence (XAI): Toward medical XAI, *IEEE Trans. Neural Netw. Learn. Syst.* 32 (11) (2020) 4793–4813.

[6] J. Yang, K. Zhou, Y. Li, Z. Liu, Generalized out-of-distribution detection: A survey, 2021, arXiv preprint arXiv:2110.11334.

[7] K. Muhammad, A. Ullah, J. Lloret, J. Del Ser, V.H.C. de Albuquerque, Deep learning for safe autonomous driving: Current challenges and future directions, *IEEE Trans. Intell. Transp. Syst.* 22 (7) (2020) 4316–4336.

[8] A. Uwimana, R. Senanayake, Out of distribution detection and adversarial attacks on deep neural networks for robust medical image analysis, 2021, arXiv preprint arXiv:2107.04882.

[9] T. Cao, C.-W. Huang, D.Y.-T. Hui, J.P. Cohen, A benchmark of medical out of distribution detection, 2020, arXiv preprint arXiv:2007.04250.

[10] F. Akopyan, J. Sawada, A. Cassidy, R. Alvarez-Icaza, J. Arthur, P. Merolla, N. Imam, Y. Nakamura, P. Datta, G.-J. Nam, B. Taba, M. Beakes, B. Brezzo, J.B. Kuang, R. Manohar, W.P. Risk, B. Jackson, D.S. Modha, TrueNorth: Design and tool flow of a 65 mW 1 million neuron programmable neuromorphic chip, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* 34 (10) (2015) 1537–1557.

[11] S. Ghosh-Dastidar, H. Adeli, Third generation neural networks: Spiking neural networks, in: *Advances in Computational Intelligence*, Springer, 2009, pp. 167–178.

[12] L. Paulun, A. Wendt, N. Kasabov, A retinotopic spiking neural network system for accurate recognition of moving objects using NeuCube and dynamic vision sensors, *Front. Comput. Neurosci.* 12 (2018) 42.

[13] C. Geng, S.-j. Huang, S. Chen, Recent advances in open set recognition: A survey, *IEEE Trans. Pattern Anal. Mach. Intell.* 43 (10) (2020) 3614–3631.

[14] D. Hendrycks, K. Gimpel, A baseline for detecting misclassified and out-of-distribution examples in neural networks, 2016, arXiv preprint arXiv:1610.02136.

[15] S. Liang, Y. Li, R. Srikant, Enhancing the reliability of out-of-distribution image detection in neural networks, in: *Proceedings of the 6th International Conference on Learning Representations, ICLR*, 2018.

[16] K. Lee, K. Lee, H. Lee, J. Shin, A simple unified framework for detecting out-of-distribution samples and adversarial attacks, *Adv. Neural Inf. Process. Syst.* 31 (2018).

[17] R. Kamoi, K. Kobayashi, Why is the mahalanobis distance effective for anomaly detection? 2020, arXiv:2003.00402.

[18] D. Hendrycks, M. Mazeika, T. Dietterich, Deep anomaly detection with outlier exposure, in: *Proceedings of the International Conference on Learning Representations, ICLR*, 2019.

[19] W. Liu, X. Wang, J. Owens, Y. Li, Energy-based out-of-distribution detection, *Adv. Neural Inf. Process. Syst.* 33 (2020) 21464–21475.

[20] Z. Lin, S.D. Roy, Y. Li, Mood: Multi-level out-of-distribution detection, in: *Proceedings of the IEEE/CVF Computer Vision and Pattern Recognition Conference, CVPR*, 2021, pp. 15313–15323.

[21] H. Wang, W. Liu, A. Bocchieri, Y. Li, Can multi-label classification networks know what they don't know? *NeurIPS* 34 (2021).

[22] R. Huang, A. Geng, Y. Li, On the importance of gradients for detecting distributional shifts in the wild, *Adv. Neural Inf. Process. Syst.* 34 (2021) 677–689.

[23] Y. Sun, C. Guo, Y. Li, ReAct: Out-of-distribution detection with rectified activations, *Adv. Neural Inf. Process. Syst.* 34 (2021) 144–157.

[24] Y. Song, N. Sebe, W. Wang, RankFeat: Rank-1 feature removal for out-of-distribution detection, in: *NeurIPS*, 2022.

[25] S.B. Furber, F. Galluppi, S. Temple, L.A. Plana, The SpiNNaker project, *Proc. IEEE* 102 (5) (2014) 652–665.

[26] M. Davies, N. Srinivasa, T.-H. Lin, G. China, Y. Cao, S.H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C.-K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataramanan, Y.-H. Weng, A. Wild, Y. Yang, H. Wang, Loihi: A neuromorphic manycore processor with on-chip learning, *IEEE Micro* 38 (1) (2018) 82–99.

[27] X. Wang, X. Lin, X. Dang, Supervised learning in spiking neural networks: A review of algorithms and evaluations, *Neural Netw.* 125 (2020) 258–280.

[28] D. Auge, J. Hille, E. Mueller, A. Knoll, A survey of encoding techniques for signal processing in spiking neural networks, *Neural Process. Lett.* 53 (6) (2021) 4693–4710.

[29] N.K. Kasabov, *Time-space, spiking neural networks and brain-inspired artificial intelligence*, Springer, 2019.

[30] W. Gerstner, W.M. Kistler, *Spiking Neuron Models: Single Neurons, Populations, Plasticity*, Cambridge University Press, 2002.

[31] E. Izhikevich, Simple model of spiking neurons, *IEEE Trans. Neural Netw.* 14 (6) (2003) 1569–1572.

[32] J.L. Lobo, J. Del Ser, A. Bifet, N. Kasabov, Spiking neural networks and online learning: An overview and perspectives, *Neural Netw.* 121 (2020) 88–100.

[33] N. Caporale, Y. Dan, Spike timing-dependent plasticity: A Hebbian learning rule, *Annu. Rev. Neurosci.* 31 (2008) 25–46.

[34] S.M. Bohte, J.N. Kok, H. La Poutre, Error-backpropagation in temporally encoded networks of spiking neurons, *Neurocomputing* 48 (1–4) (2002) 17–37.

[35] O. Booi, H. tat Nguyen, A gradient descent rule for spiking neurons emitting multiple spikes, *Inform. Process. Lett.* 95 (6) (2005) 552–558.

- [36] A. Mohammed, S. Schliebs, S. Matsuda, N. Kasabov, Span: Spike pattern association neuron for learning spatio-temporal spike patterns, *Int. J. Neural Syst.* 22 (04) (2012) 1250012.
- [37] E.O. Neftci, H. Mostafa, F. Zenke, Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks, *IEEE Signal Process. Mag.* 36 (6) (2019) 51–63.
- [38] F. Zenke, S. Ganguli, Superspike: Supervised learning in multilayer spiking neural networks, *Neural Comput.* 30 (6) (2018) 1514–1541.
- [39] A. Jeyasothy, S. Sundaram, S. Ramasamy, N. Sundararajan, A novel method for extracting interpretable knowledge from a spiking neural classifier with time-varying synaptic weights, 2019, arXiv preprint arXiv:1904.11367.
- [40] A. Jeyasothy, S. Sundaram, N. Sundararajan, SEFRON: A new spiking neuron model with time-varying synaptic efficacy function for pattern classification, *IEEE Trans. Neural Netw. Learn. Syst.* 30 (4) (2018) 1231–1240.
- [41] S. Soltic, N. Kasabov, Knowledge extraction from evolving spiking neural networks with rank order population coding, *Int. J. Neural Syst.* 20 (06) (2010) 437–445.
- [42] M. Dobarjeh, Z. Dobarjeh, N. Kasabov, M. Barati, G.Y. Wang, Deep learning of explainable EEG patterns as dynamic spatiotemporal clusters and rules in a Brain-Inspired spiking neural network, *Sensors* 21 (14) (2021) 4900.
- [43] K. Kumarasinghe, N. Kasabov, D. Taylor, Deep learning and deep knowledge representation in spiking neural networks for brain-computer interfaces, *Neural Netw.* 121 (2020) 169–185.
- [44] N.K. Kasabov, NeuCube: A spiking neural network architecture for mapping, learning and understanding of spatio-temporal brain data, *Neural Netw.* 52 (2014) 62–76.
- [45] Y. Kim, P. Panda, Visual explanations from spiking neural networks using inter-spike intervals, *Sci. Rep.* 11 (1) (2021) 1–14.
- [46] E. Nguyen, Temporal Spike Attribution: A Local Feature-Based Explanation for Temporally Coded Spiking Neural Networks, University of Twente, 2021.
- [47] Y. Wu, L. Deng, G. Li, J. Zhu, Y. Xie, L. Shi, Direct training for spiking neural networks: Faster, larger, better, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, no. 01, 2019, pp. 1311–1318.
- [48] A. Samadzadeh, F.S.T. Far, A. Javadi, A. Nickabadi, M.H. Chehreghani, Convolutional spiking neural networks for spatio-temporal feature extraction, 2020, arXiv preprint arXiv:2003.12346.
- [49] S. Lian, Q. Liu, R. Yan, G. Pan, H. Tang, Training deep convolutional spiking neural networks with spike probabilistic global pooling, *Neural Comput.* 34 (5) (2022) 1170–1188.
- [50] R. Vaila, Deep Convolutional Spiking Neural Networks for Image Classification, Boise State University, 2021.
- [51] D. Heeger, et al., Poisson model of spike generation, Vol. 5, no. 1–13, Handout, University of Stanford, 2000, p. 76.
- [52] C.C. Aggarwal, A. Hinneburg, D.A. Keim, On the surprising behavior of distance metrics in high dimensional space, in: *International Conference on Database Theory*, Springer, 2001, pp. 420–434.
- [53] A.M. Garcia-Vico, F. Herrera, A preliminary analysis on software frameworks for the development of spiking neural networks, in: *International Conference on Hybrid Artificial Intelligence Systems*, HAIS, Springer, 2021, pp. 564–575.
- [54] W. Gerstner, W.M. Kistler, R. Naud, L. Paninski, *Neuronal dynamics: From single neurons to networks and models of cognition*, Cambridge University Press, 2014.
- [55] C. Pehle, J.E. Pedersen, Norse-A Deep Learning Library for Spiking Neural Networks, Vol. 4422025, Zenodo, 2021, Documentation: <https://norse.ai/docs/> [Online]. Available: <https://doi.org/10.5281/zenodo>.
- [56] L. Deng, The mnist database of handwritten digit images for machine learning research, *IEEE Signal Process. Mag.* 29 (6) (2012) 141–142.
- [57] H. Xiao, K. Rasul, R. Vollgraf, Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017, arXiv preprint arXiv:1708.07747.
- [58] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, D. Ha, Deep learning for classical japanese literature, 2018, arXiv preprint arXiv:1812.01718.
- [59] G. Cohen, S. Afshar, J. Tapson, A. Van Schaik, EMNIST: Extending MNIST to handwritten letters, in: *2017 International Joint Conference on Neural Networks, IJCNN, IEEE*, 2017, pp. 2921–2926.
- [60] G. Orchard, A. Jayawant, G.K. Cohen, N. Thakor, Converting static image datasets to spiking neuromorphic datasets using saccades, *Front. Neurosci.* 9 (2015) 437.
- [61] A. Amir, B. Taba, D. Berg, T. Melano, J. McKinstry, C. Di Nolfo, T. Nayak, A. Andreopoulos, G. Garreau, M. Mendoza, et al., A low power, fully event-based gesture recognition system, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7243–7252.
- [62] Y. Bulatov, Notmnist dataset, Vol. 2, Tech. Rep., Google (Books/OCR), 2011, [Online] Available: <http://yaroslavvb.blogspot.it/2011/09/notmnist-dataset.html>.
- [63] B.M. Lake, R. Salakhutdinov, J.B. Tenenbaum, Human-level concept learning through probabilistic program induction, *Science* 350 (6266) (2015) 1332–1338.
- [64] A. Krizhevsky, G. Hinton, et al., Learning Multiple Layers of Features from Tiny Images, Citeseer, 2009.
- [65] H. Li, H. Liu, X. Ji, G. Li, L. Shi, Cifar10-dvs: An event-stream dataset for object classification, *Front. Neurosci.* 11 (2017) 309.
- [66] G. Tan, Y. Wang, H. Han, Y. Cao, F. Wu, Z.-J. Zha, Multi-grained spatio-temporal features perceived network for event-based lip-reading, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 20094–20103.
- [67] N. Mu, J. Gilmer, MNIST-c: A robustness benchmark for computer vision, 2019, arXiv preprint arXiv:1906.02337.
- [68] J. Demšar, Statistical comparisons of classifiers over multiple data sets, *J. Mach. Learn. Res.* 7 (2006) 1–30.
- [69] A. Benavoli, G. Corani, J. Demšar, M. Zaffalon, Time for a change: A tutorial for comparing multiple classifiers through Bayesian analysis, *J. Mach. Learn. Res.* 18 (1) (2017) 2653–2688.
- [70] R.S. Johansson, I. Birznieks, First spikes in ensembles of human tactile afferents code complex spatial fingertip events, *Nature Neurosci.* 7 (2) (2004) 170–177.
- [71] N.K. Kasabov, Spiking neural networks for deep learning and knowledge representation: Editorial, *Neural Netw.* 119 (2019) 341–342.
- [72] I. Abouhassan, N. Kasabov, G. Popov, R. Trifonov, Why use evolving neuro-fuzzy and spiking neural networks for incremental and explainable learning of time series? A case study on predictive modelling of trade imports and outlier detection, in: *IEEE 11th International Conference on Intelligent Systems*, IS, 2022, pp. 1–7.