Abstraction and Symbolic Execution of Deep Neural Networks with Bayesian Approximation of Hidden Features

Nicolas Berthier[1], Amany Alshareef[1], James Sharp[2],
Sven Schewe[1], Xiaowei Huang[1]

[1] University of Liverpool, UK
[2] Defence Science and Technology Laboratory (Dstl), UK

Abstract

Intensive research has been conducted on the verification and validation of deep neural networks (DNNs), aiming to understand if, and how, DNNs can be applied to safety critical applications. However, existing verification and validation techniques are limited by their scalability, over both the size of the DNN and the size of the dataset. In this paper, we propose a novel abstraction method which abstracts a DNN and a dataset into a Bayesian network (BN). We make use of dimensionality reduction techniques to identify hidden features that have been learned by hidden layers of the DNN, and associate each hidden feature with a node of the BN. On this BN, we can conduct probabilistic inference to understand the behaviours of the DNN processing data. More importantly, we can derive a runtime monitoring approach to detect in operational time rare inputs and covariate shift of the input data. We can also adapt existing structural coverage-guided testing techniques (*i.e.,* based on low-level elements of the DNN such as neurons), in order to generate test cases that better exercise hidden features. We implement and evaluate the BN abstraction technique using our DeepConcolic tool available at https://github.com/TrustAI/DeepConcolic.

1 Introduction

Neural networks generally work with high precision, but recent work has shown that they are subject to weaknesses such as adversarial attacks [35], data poisoning attacks [2], Trojan attacks [22], model inversion attacks [8], etc. Given their importance and such weaknesses, the analysis of DNNs has become a popular research direction, with research on formal verification, coverage-guided testing, etc. See [16] for a recent survey. The large size of DNNs, containing tens of thousands of neurons that interact with each other in intricate ways, leads to the scalability problem of these analysis methods, particularly for white-box

analysis techniques, whose computational complexity is usually measured over either the number of neurons or the number of parameters. Moreover, most analysis methods – either verification or testing – work with local inputs individually, which easily leads to the other scalability problem when they need to work with a large dataset. The above scalability problems, due to the size of DNNs and the size of datasets, imply the need for an abstraction technique that approximates a DNN and a dataset into a simpler model – while preserving critical properties – to improve the scalability of the analysis methods.

Contribution. This paper proposes a novel abstraction technique for DNNs through Bayesian approximation: we abstract the behaviour of a DNN on a dataset into a Bayesian network (BN), which is a probabilistic graphical model based on high-level features instead of low-level neurons.

The first step of this abstraction is naturally to identify the relevant *hidden features* for each hidden layer. Standard feature extraction techniques are available, and while we obviously have to select some for experimental results and explore a few different ways of conducting feature extraction, we predominantly treat feature extraction as black-box techniques.

The features not only provide the structure of the BN, they also provide, for each input from a given dataset[1], an observation of occurrences of features in neighbouring layers.

This first allows us to conduct probabilistic reasoning to understand how the appearance of some input features may statistically affect the appearance of hidden features or output labels, and how the appearance of some output label may affect the appearance probability of some input feature, etc. Such reasoning contributes to a global explainable AI method, helping the users to understand how the DNN behaves when processing inputs.

Second, we can use the BN as the basis for a simple monitor, which can be used to identify cases where the classification of the DNN is unreliable. In particular, we want to identify two situations where the classification of the DNN becomes unreliable: outliers (data that has hardly ever, or not at all, occurred during training), and a significantly different long-term behaviour compared to the training data-based abstraction. Such a significant **covariate shift** suggests that the network is used in an environment, which is significantly different from its training environment, and therefore its decision making might not be reliable. An **outlier** suggests that the decision making from the network might not be reliable. Both cases may **compromise the safe application of the DNN**.

Third, we design a few coverage metrics over the BN, together with a novel concolic testing method – based on symbolic execution – considering the information from both the network and the BN. These metrics and methods improve over earlier structural coverage metrics [33] and concolic testing method [32] for

---

[1]For most methods, the dataset that has been used to learning a DNN is also used to identify the features. But we also use fixed features, e.g. when a DNN is used over a new dataset, and we are also keen to keep this part of the analysis independent from how the features are selected.

DNNs. Since the abstract model can be significantly smaller than the original DNN, the effectiveness of the symbolic execution technique for generating relevant new test cases for the DNN is significantly improved.

In addition to the improved scalability, we show that this approach is closer to the design intent for deep learning. In convolutional neural networks, while neurons are the minimal (syntactic) units, the features on either the input or the hidden layers are the minimal **semantic** units. As such, the main aim of both the DNN architecture, and the back-propagation learning algorithm, is to improve the *hidden features* learned by hidden layers. There also exist significant research directions, such as interpretability, that aim at understanding what has been learned by a DNN, *e.g.,* by visualising the hidden features. We therefore base our work on the key hypothesis of deep learning that **hidden features should be first-class citizens**, rather than neurons, who do not represent a clear semantics *per se*.

The organisation of the paper is as follows. In the next section, we will present preliminaries about deep neural networks, dimensionality reduction, and Bayesian Networks. In Section 3, we will turn to the presentation of our Bayesian Network-based approach, including a generic framework on extracting hidden features for given layers (Section 3.1) and feature space discretisation (Section 3.2), concrete strategies for feature extraction and discretisation (Section 3.3), and the Bayesian network construction (Section 3.4). We then turn in Section 4 to our approach for exploiting the Bayesian Network for the detection of unanticipated operational data. Our Bayesian Network-based approach is further utilised in Section 5 for defining Bayesian Network-based hidden feature coverage metrics. We discuss in Section 6 our adaptation of a concolic testing algorithm to achieve the latter coverage, and present in Section 7 our implementation and evaluation. Finally, we will discuss related works in Section 8 and conclude in Section 9.

2 Preliminaries

Deep Neural Networks   Let $\mathcal{N}$ be a deep neural network (DNN) of a given architecture. For a learning model, we use (X, Y) to denote the training data, where X is a vector of inputs and Y is a corresponding vector of outputs such that $|X| = |Y|$. Let $\mathbb{D}_X$ be the input domain and $\mathbb{D}_Y$ be the set of labels[2]. Hence, $X \subset \mathbb{D}_X$. We may use x and y to range over $\mathbb{D}_X$ and $\mathbb{D}_Y$, respectively.

A network $\mathcal{N} \colon \mathbb{D}_X \to \mathcal{D}(\mathbb{D}_Y)$ can be seen as a mapping from $\mathbb{D}_X$ to probabilistic distributions over $\mathbb{D}_Y$. This means that $\mathcal{N}(x)$ is a probabilistic distribution, which associates each possible label $y \in \mathbb{D}_Y$ with a probability value (or confidence level) $(\mathcal{N}(x))_y$. We consequently let $f_\mathcal{N} \colon \mathbb{D}_X \to \mathbb{D}_Y$ and $c_\mathcal{N} \colon \mathbb{D}_X \to [0, 1]$ be such that, for any $x \in \mathbb{D}_X$,

$$f_\mathcal{N}(x) = \arg\max_{y \in \mathbb{D}_Y}\{(\mathcal{N}(x))_y\}.$$

---

[2]Throughout this paper, we will use double-struck capitals to denote domains.

$f_\mathcal{N}(x)$ returns the label with the greatest probability, *i.e.,* the classification label.

In this work, we consider *sequential* networks: a network $\mathcal{N}$ consists of a sequence of layers $\mathsf{Layers} = (l_1, \ldots, l_K)$, where every layer $l_i \in \mathsf{Layers}$ contains a set of $|l_i|$ *neurons* $h_i = \{n_{i,1}, \ldots, n_{i,|l_i|}\}$. Each neuron $n_{i,j}$ in h_i first computes a value $\hat{n}_{i,j}$ in some domain that is usually assumed to be the set of Real numbers $\mathbb{R}$, and this value is typically computed as a function of the outputs of neurons in layer l_{i-1}. The particular semantics of the latter function depends on the role of the layer: this function often consists of a *linear combination* of (a portion of) its inputs, which means that one has

$$\hat{n}_{i,j} = W_{i,j}\left(n_{i-1,1}, \ldots, n_{i-1,|l_{i-1}|}\right) + b_i \tag{1}$$

where $W_{i,j}$ and b_i are coefficients of *weight* and *bias* matrices for the layer, that are the parameters learned by the network during the training phase.

Other functions are also useful, notably to reduce overfitting by decreasing the capacity of the network (*i.e.,* the total number of weight and bias parameters that need to be learned). This is for instance the case of *max-pooling* layers, that down-sample their inputs by selecting the maximum over a given subset:

$$\hat{n}_{i,j} = \max s_j\left(n_{i-1,1}, \ldots, n_{i-1,|l_{i-1}|}\right) \tag{2}$$

where each s_j, for $j \in \{1, \ldots |l_i|\}$, is a function that gives a subset of all given neuron outputs from the preceding layer (the successive j's typically result in a sliding window over all inputs). The functional semantics of such a layer is therefore non-linear.

Each value $\hat{n}_{i,j}$ of a layer l_i may additionally pass through an *activation function* σ_i that enables $\mathcal{N}$ to capture non-linearities: the output of neuron $n_{i,j}$ is called its *activation*, and is then computed as

$$n_{i,j} = \sigma_i(\hat{n}_{i,j}). \tag{3}$$

Notable examples of activation functions include *e.g.,* ReLU, Sigmoid, tanh, etc. In this work, we shall focus on the former, which can be defined as:

$$\mathrm{ReLU}(\hat{n}) \stackrel{\mathrm{def}}{=} \max\{0, \hat{n}\} \tag{4}$$

The product of domain of all neuron values in h_i builds up the *valuation space* $\mathbb{L}_i$ of all neurons for layer l_i (the $\hat{n}_{i,j}$'s); one typically has, therefore, $\mathbb{L}_i = \mathbb{R}^{|l_i|}$. We pose that the valuation of every neuron in the *input layer* l_1 corresponds to each component (*e.g.,* pixel) of the input $x \in \mathbb{D}_X$, so that $\mathbb{L}_1 = \mathbb{D}_X$. We also denote with $\hat{h}_i(x) \in \mathbb{L}_i$ the set of all *neuron values* in layer l_i (before being passed through any subsequent activation function), when $\mathcal{N}$ is fed with input x.

Figure 2 presents a simple four layer network, where each of the first three layers contains 4 neurons and the last layer contains 2 neurons.

Layer	Function specification	Output shape	#weights + #bias parameters
l_0	conv2d (convolutional)	$26 \times 26 \times 8$	80
l_1	flatten (flat)	5408	0
l_2	dense (dense)	42	227 178
l_3	dense_1 (dense)	10	430

Table 1: Structure of the DNN $\mathcal{N}_{\mathsf{ms}}$, dedicated to the MNIST dataset; layers l_0 and l_2 incorporate ReLU activation functions, whereas the output layer l_3 involves a classical Softmax to output the predicted label.

Layer	Function specification	Output shape	#weights + #bias parameters
l_0	conv2d (convolutional)	$26 \times 26 \times 8$	80
l_1	max_pooling2d (max-pooling)	$13 \times 13 \times 8$	0
l_1	flatten (flat)	1352	0
l_2	dense (dense)	42	56 826
l_3	dense_1 (dense)	10	430

Table 2: Structure of the DNN $\mathcal{N}_{\mathsf{mx}}$, similar to $\mathcal{N}_{\mathsf{ms}}$ except for it incorporates a max-pooling layer.

Example DNNs for Handwritten Digits Recognition For the purpose of keeping our illustrations simple, we base most of our examples on a small DNN $\mathcal{N}_{\mathsf{ms}}$, whose structure is summarised in Table 1. We have trained $\mathcal{N}_{\mathsf{ms}}$ on the MNIST dataset, that consists of labelled black-and-white images of 28×28 pixels, until it reached a validation accuracy of about 98%. We have similarly constructed and trained a DNN $\mathcal{N}_{\mathsf{mx}}$ that incorporates a max-pooling layer, summarised in Table 2, to further investigate the effect of such filters on both the hidden space and the ability of our test generation algorithm to capture their intrinsically non-linear functional semantics.

Dimensionality Reduction via Feature Extraction The goal of *dimensionality reduction* techniques consists in computing a mapping from a high-dimensional space into some space of (much) lower dimensionality, called the *feature space*. Computing approaches for such mappings usually rely on statistical principles and operate on a given sample of high-dimensional data; prominent examples include *Principal Component Analysis* (PCA), *t-distributed Stochastic Neighbour Embedding* (t-SNE) [25], or other forms of *kernel tricks* [36]. Other techniques that do not specifically address dimensionality reduction *per se*, still compute relevant transformations in our context—such is the case of *Independent Component Analysis* (ICA) [17]. Each one of these techniques results in mappings that satisfy various properties of interest, among which we notice:

Linearity: this property is satisfied when the obtained mapping consists of a pair of matrices $W \in \mathbb{D}^{H \times L}$ and $B \in \mathbb{D}^L$ for some field $\mathbb{D}$, such that any element $x \in \mathbb{D}^H$ from the H-dimensional space is mapped onto the element $xW + B \in \mathbb{D}^L$ in the L-dimensional feature space. Both PCA and ICA belong to the set of techniques that give such linear mappings.

Orthogonality: PCA computes features one by one, in an effort to minimise the variance of the data. The importance of the features – in terms of

their effectiveness in the reduction of variance – is gradually decreased, although the features are orthogonal (linear independence) to each other. Conversely, ICA computes features at the same time, and the generated features may not be orthogonal to each other.

Many available dimensionality reduction techniques compute mappings that are non-linear, such as t-SNE for instance, which is a supervised learning procedure that tries to separate clusters of samples of the same label for adequate visual representations in 2D or 3D spaces. Whereas some of these techniques are adequate for constructing our BN abstraction, we focus in this work on feature mappings that are *linear*. This restriction will actually enable us to provide an effective feature-based test-case generation algorithm.

We will consider feature mappings in decomposed form: a mapping will be given as a set $\Lambda = \{\lambda_j\}_{j \in \{1,\ldots,|\Lambda|\}}$, where each λ_j maps the high-dimensional space into the jth-component of the feature space. We will use a double-struck capital $\mathbb{F}$ to denote such a feature space, and indexed $\mathbb{F}$'s to denote the domain of each one of its individual components (typically $\mathbb{R}$).

Bayesian Network A *Bayesian Network* (BN) $\mathcal{B} = (V, E, P)$ is a directed acyclic graph (DAG) whose nodes V represent variables in the Bayesian sense [20], E are edges, and P maps each node in V to a set of *probability tables*. $\mathcal{B}$ is such that each edge in E represents conditional dependencies, and nodes that are not connected (no path connects one node to another) represent variables that are conditionally independent of each other. In our BN-based analysis, every node that has at least one incoming edge will be attached with a *conditional probability table*, and each source node will be associated with a *marginal probability table*.

Figure 1 presents a simple BN abstracted from $\mathcal{N}_{\mathsf{ms}}$ according to the process described in the next Section. It represents 6 binary variables (that take their values in the set $\{0, 1\}$), and it thus has 6 nodes, each of which is attached with a probability table. The two nodes conv2d.0 and conv2d.1 have no predecessor, and they are thus only associated with marginal probability tables. From the first table, we can learn that the variable conv2d.0 takes value 0 with probability 0.5. Every other table in this Figure represents a set of conditional probabilities: for instance, under the condition that variables dense.0 and dense.1 respectively hold values 1 and 0, then the variable dense_1.0 takes value 0 with probability 0.05.

3 Bayesian Network-based Abstraction of DNN

Let us now turn to the description of our approach for constructing a Bayesian Network in a way that suits out needs for capturing the distribution of neuron valuations in terms of hidden features encoded in each DNN layers, as well as their causal relationships.

conv2d.0		conv2d.1	
0	0.5	0	0.50
1	0.5	1	0.50

conv2d.0	conv2d.1	dense.0	
0	0	0	0.74
0	0	1	0.26
0	1	0	0.92
0	1	1	0.08
1	0	0	0.11
1	0	1	0.89
1	1	0	0.23
1	1	1	0.76

conv2d.0	conv2d.1	dense.1	
0	0	0	0.72
0	0	1	0.28
0	1	0	0.54
0	1	1	0.46
1	0	0	0.38
1	0	1	0.61
1	1	0	0.35
1	1	1	0.65

conv2d.0   conv2d.1

dense.0   dense.1

dense_1.0   dense_1.1

dense.0	dense.1	dense_1.0	
0	0	0	0.55
0	0	1	0.45
0	1	0	0.99
0	1	1	0.01
1	0	0	0.05
1	0	1	0.95
1	1	0	0.42
1	1	1	0.58

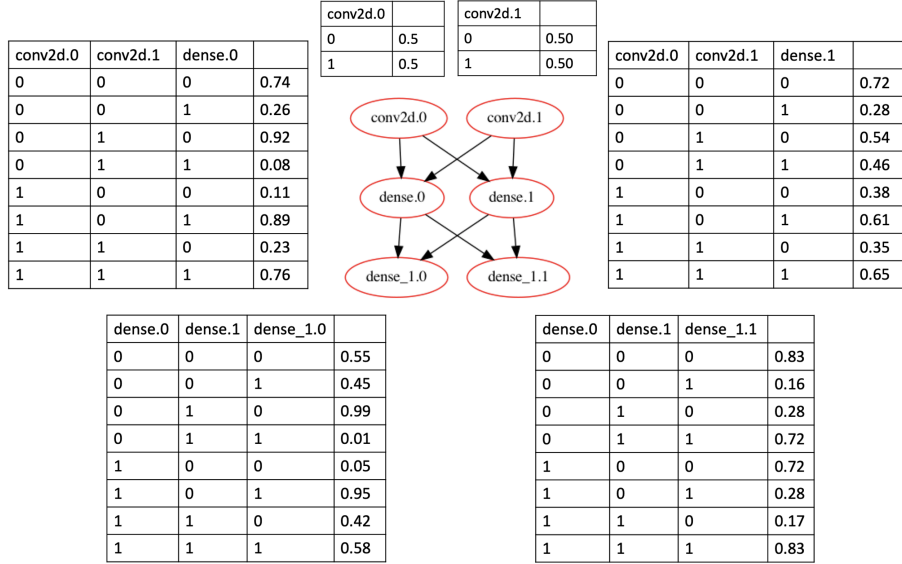| dense.0	dense.1	dense_1.1	
0	0	0	0.83
0	0	1	0.16
0	1	0	0.28
0	1	1	0.72
1	0	0	0.72
1	0	1	0.28
1	1	0	0.17
1	1	1	0.83

Figure 1: An example Bayesian network abstracted from a network of three hidden layers (one convolutional followed by two fully-connected layers) trained on MNIST dataset.

3.1   Extraction of Hidden Features

Let us assume that some feature extraction technique from Section 2 has been used to analyse the set of neuron values $\hat{n}_{i,1}, \ldots, \hat{n}_{i,|l_i|}$ that are induced by all inputs in a given *training set* $X_{train} \subset \mathbb{D}_X$ at a given layer $l_i \in$ Layers. This produces a set of feature mappings $\Lambda_i = \{\lambda_{i,j}\}_{j \in \{1,\ldots,|\Lambda_i|\}}$, where each $\lambda_{i,j} \colon \mathbb{L}_i \to \mathbb{F}_{i,j}$ maps the neuron valuation space $\mathbb{L}_i$ into the j-th component of the feature space $\mathbb{F}_i$ for layer l_i. The latter space is the product $\mathbb{F}_i \stackrel{\text{def}}{=} \prod_{j \in \{1,\ldots,|\Lambda_i|\}} \mathbb{F}_{i,j}$.

Further, Λ_i is such that the neuron values $\hat{h}_i(x)$ – *i.e.,* before being fed to the activation function σ_i – for any input $x \in \mathbb{D}_X$, can be transformed into a $|\Lambda_i|$-dimensional vector

$$\langle \lambda_{i,1} \circ \hat{h}_i(x), \ldots, \lambda_{i,|\Lambda_i|} \circ \hat{h}_i(x) \rangle \in \mathbb{F}_i \tag{5}$$

where $\lambda_{i,j} \circ \hat{h}_i(x)$ represents the j-th component of the value obtained after mapping $\hat{h}_i(x)$ into the feature space. We will refer to the projection $\lambda_{i,j} \circ \hat{h}_i(x)$ as the *hidden feature valuation induced by x on component $\mathbb{F}_{i,j}$*.

Figure 2 gives an illustrative diagram of reducing h_1, h_2, h_3 to features. In particular, each h_i is reduced to two features $\lambda_{i,1} \circ h_i$ and $\lambda_{i,2} \circ h_i$.
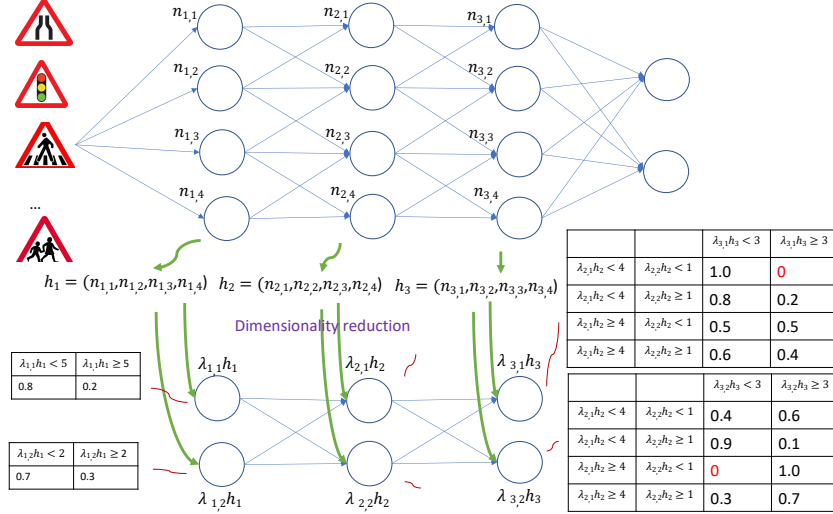
		$\lambda_{3,1}h_3 < 3$	$\lambda_{3,1}h_3 \geq 3$
$\lambda_{2,1}h_2 < 4$	$\lambda_{2,2}h_2 < 1$	1.0	0
$\lambda_{2,1}h_2 < 4$	$\lambda_{2,2}h_2 \geq 1$	0.8	0.2
$\lambda_{2,1}h_2 \geq 4$	$\lambda_{2,2}h_2 < 1$	0.5	0.5
$\lambda_{2,1}h_2 \geq 4$	$\lambda_{2,2}h_2 \geq 1$	0.6	0.4

		$\lambda_{3,2}h_3 < 3$	$\lambda_{3,2}h_3 \geq 3$
$\lambda_{2,1}h_2 < 4$	$\lambda_{2,2}h_2 < 1$	0.4	0.6
$\lambda_{2,1}h_2 < 4$	$\lambda_{2,2}h_2 \geq 1$	0.9	0.1
$\lambda_{2,1}h_2 \geq 4$	$\lambda_{2,2}h_2 < 1$	0	1.0
$\lambda_{2,1}h_2 \geq 4$	$\lambda_{2,2}h_2 \geq 1$	0.3	0.7

$\lambda_{1,1}h_1 < 5$	$\lambda_{1,1}h_1 \geq 5$
0.8	0.2

$\lambda_{1,2}h_1 < 2$	$\lambda_{1,2}h_1 \geq 2$
0.7	0.3

$h_1 = (n_{1,1}, n_{1,2}, n_{1,3}, n_{1,4})$ $h_2 = (n_{2,1}, n_{2,2}, n_{2,3}, n_{2,4})$ $h_3 = (n_{3,1}, n_{3,2}, n_{3,3}, n_{3,4})$

Dimensionality reduction

Figure 2: Illustration of our Reduction of Neural Networks to Bayesian Networks

3.2 Discretisation of the Hidden Feature Space

The feature extraction techniques mentioned in Section 2 result in mappings $\lambda_{i,j}$ that range over a continuous and potentially infinite domain, such as $\mathbb{R}$. Yet, our BN-based abstraction technique relies on the construction of *probability tables*, where each entry associates a set of *distinct* hidden feature values with a probability. For this construction to be relevant, we therefore *discretise* each hidden feature component into a *finite* set of sub-spaces. We first introduce the general process and notations we use for discretising hidden features, and elaborate on various strategies in the next Section.

Discretisation Process   We define the discretisation process in a similar way as the feature extraction above, and rely on a set of *training inputs* X_{train} to gather statistics about the projections onto each hidden feature component of the resulting neuron values. Given such a component $\mathbb{F}$, we use these aggregate statistics to find a *partition* of $\mathbb{F}$.

We give in Figure 3 an illustration of what the discretisation process produces. Each plot focuses on one hidden feature component $\mathbb{F}_{\mathtt{dense},j}$ of a layer $\mathtt{dense}$, and the clouds of dots represents the respective distribution of hidden feature values $\lambda_{\mathtt{dense},j} \circ \hat{h}_{\mathtt{dense}}(x)$, for x uniformly drawn from a subset of X_{train}. The goal of the discretisation process is therefore to partition each horizontal axis into a set of distinct regions. In this work, we only consider partitions into finite sets of intervals. We describe the remaining elements of Figure 3 in the next Section.

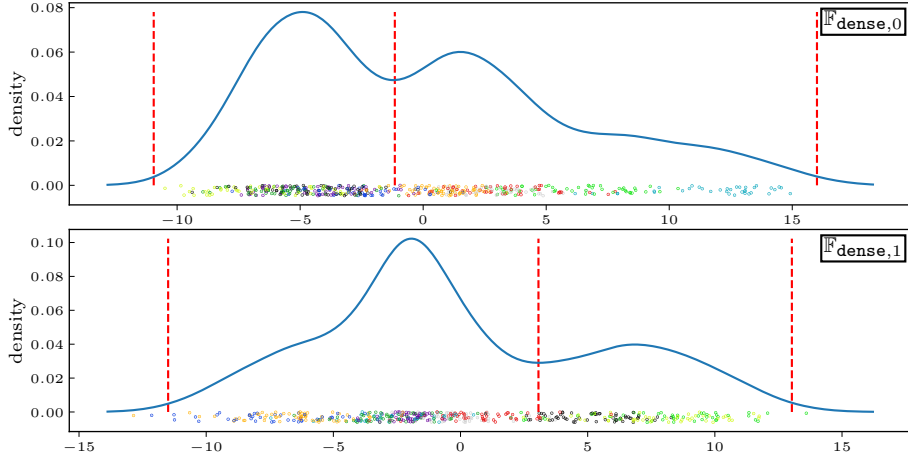More specifically, the discretisation process produces two objects for each mapping $\lambda_{i,j}$:

8

Figure 3: Projection onto two hidden feature components $\mathbb{F}_{\texttt{dense},0}$ and $\mathbb{F}_{\texttt{dense},1}$ of neuron values induced by a sample of training data X_{train}, associated density estimates (solid lines), and interval boundaries for discretisation (dashed vertical lines). Although our current strategies for discretisation ignore input labels, we still represent them by colouring the respective dots.

- a *finite set of intervals* that partitions the feature component;
- a total *discretisation function* that maps each neuron valuation for this layer, *i.e.,* that belongs to $\mathbb{L}_i$, onto the corresponding interval.

Further Notations   We use $\sharp$ exponents to denote discrete spaces, or entities that belong to a discrete space. Formally, the process of discretising a feature mapping $\lambda_{i,j}$ consists in finding a finite set of m right-open intervals $\mathbb{F}_{i,j}^{\sharp} = \left\{ f_{i,j}^{\sharp 1}, \ldots, f_{i,j}^{\sharp m} \right\}$ that partitions the codomain $\mathbb{F}_{i,j}$ of $\lambda_{i,j}$.

Example 3.1 (Lidden feature partitioning)**.** We may partition $\mathbb{F}_{i,j} = (-\infty, \infty)$ into two intervals $\mathbb{F}_{i,j}^{\sharp} = \left\{ f_{i,j}^{\sharp 1}, f_{i,j}^{\sharp 2} \right\}$ such that $f_{i,j}^{\sharp 1}$ includes all real numbers strictly lower than 5, and $f_{i,j}^{\sharp 2}$ includes all real numbers greater or equal than 5.

Given an interval $f_{i,j}^{\sharp k}$, we denote with $f_{i,j}^{[k}$ (resp. $f_{i,j}^{k[}$) the least element in $f_{i,j}^{\sharp k}$ (resp. the least element *not* in $f_{i,j}^{\sharp k}$). This discretisation being a partition, one always has $f_{i,j}^{[1} = -\infty$ and $f_{i,j}^{n[} = \infty$ if the component ranges over $\mathbb{R}$.

Example 3.2 (Continuing Example 3.1)**.** We can write $f_{i,j}^{\sharp 1}$ as $(-\infty, 5[$ and $f_{i,j}^{\sharp 2}$ as $[5, +\infty)$.

To obtain the hidden feature interval that results from all neuron values at layer l_i, we define the discretisation function $\mathsf{Discr}_{i,j}^{\sharp} : \mathbb{L}_i \to \mathbb{F}_{i,j}^{\sharp}$ as the total mapping that gives the interval for the jth component at layer l_i, which corresponds
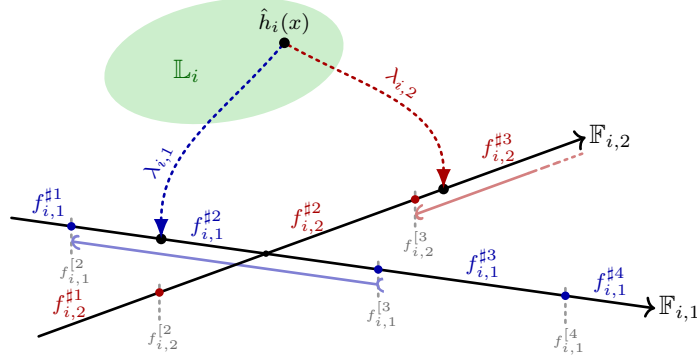
9

Figure 4: Schematic representation of the discretised feature space for a single layer l_i, and two extracted features. The ellipse represents the high-dimensional *valuation space* for l_i, *i.e.,* $\mathbb{L}_i$, and we illustrate the discretisation of two feature components $\mathbb{F}_{i,1}$ and $\mathbb{F}_{i,1}$ that are induced by mappings $\lambda_{i,1}$ and $\lambda_{i,2}$, respectively. We arbitrarily partition feature $\lambda_{i,1}$ into a set of four intervals $\mathbb{F}^\sharp_{i,1} = \left\{ f^{\sharp 1}_{i,1}, \ldots, f^{\sharp 4}_{i,1} \right\}$, and $\lambda_{i,2}$ into three intervals $\mathbb{F}^\sharp_{i,2} = \left\{ f^{\sharp 1}_{i,2}, \ldots, f^{\sharp 3}_{i,2} \right\}$. The application of $\lambda_{i,1}$ and $\lambda_{i,2}$ on the valuation $\hat{h}_i(x)$ induced by some input $x \in \mathbb{D}_X$ gives feature values $\lambda_{i,1} \circ \hat{h}_i(x)$ and $\lambda_{i,2} \circ \hat{h}_i(x)$ that belong to intervals $f^{\sharp 2}_{i,1}$ and $f^{\sharp 3}_{i,2}$, respectively.

to the interval obtained for the discretised hidden feature component $\mathbb{F}_{i,j}$; *i.e.,* $\mathsf{Discr}^\sharp_{i,j} \circ \hat{h}_i(x) = f^{\sharp k}_{i,j}$ iff $\lambda_{i,j} \circ \hat{h}_i(x) \in f^{\sharp k}_{i,j}$.

At last, we define the full *discretised hidden feature space* for a layer l_i as the product $\mathbb{F}^\sharp_i \overset{\text{def}}{=} \prod_{j \in \{1, \ldots, |\Lambda_i|\}} \mathbb{F}^\sharp_{i,j}$, and use $F^\sharp_i$ to denote elements of $\mathbb{F}^\sharp_i$.

Illustrations We illustrate in Figure 4 the notations that we use to denote the discretisation for a single layer according to some feature mapping $\Lambda_i \supseteq \{\lambda_{i,1}, \lambda_{i,2}\}$ and arbitrary discretisations to four and three intervals, respectively: *i.e.,* $\mathbb{F}^\sharp_{i,1} = \left\{ f^{\sharp 1}_{i,1}, \ldots, f^{\sharp 4}_{i,1} \right\}$ and $\mathbb{F}^\sharp_{i,2} = \left\{ f^{\sharp 1}_{i,2}, \ldots, f^{\sharp 3}_{i,2} \right\}$. We also show the respective intervals obtained from the neuron values induced by some input x as $\lambda_{i,1} \circ \hat{h}_i(x) \in f^{\sharp 2}_{i,1}$ and $\lambda_{1,2} \circ \hat{h}_i(x) \in f^{\sharp 3}_{i,2}$; in other words: the application of the feature mapping $\lambda_{i,1}$ to the neuron values $\hat{h}_i(x)$ at layer l_i obtained with input x, belongs to the second interval $f^{\sharp 2}_{i,1}$ in the discretisation of $\mathbb{F}_{i,1}$. We apply the same principle and obtain the third interval $f^{\sharp 3}_{i,2}$ for $\mathbb{F}_{i,2}$.

Example 3.3. We further give in Figure 2 a more concrete example, where the feature $\lambda_{1,1}$ is split into the set of ranges $\{(-\infty, 5[, \ [5, +\infty)\}$, and the feature component $\lambda_{2,2}$ is partitioned into $\{(-\infty, 1[, \ [1, +\infty)\}$.

10

3.3 Strategies for Hidden Feature Extraction & Discretisation

Predefined strategies need to be employed for both extracting hidden features (*i.e.*, to identify a number of mappings that correspond to relevant hidden features), and constructing a discretisation as defined above, for each of them.

At this stage, we can anticipate that the choice of such strategies has a great impact on various properties of the resulting discretised hidden feature space, such as its ability to accurately separate the hidden feature valuations that are induced by "sufficiently" dissimilar inputs, or to isolate hidden sub-spaces that correspond to corner cases. In particular, we hypothesise that various characteristics of each layer within the neural network, such as its function and position (depth), need to be taken into consideration when selecting discretisation strategies.

Furthermore, the fact that we use the training data to perform extractions and discretisations has an impact on the semantics of the resulting hidden sub-spaces. Indeed, in doing so we rely on the training data and the set of neuron valuations it induces within the DNN to infer an *abstract representation of its learned behaviours*. For now, we devise and explain several such strategies.

Density-based Discretisation   Let us now refer to the density estimates given in Figure 3. These estimations were obtained using Kernel Density Estimation (with Gaussian kernels), which is a standard technique for empirically approximating a probability distribution over a given domain (here, each $\mathbb{F}_{\mathtt{dense},j}$). We can first observe that the distributions of hidden feature values exhibit some irregularities that can be exploited to identify relevant intervals.

We adapt DeepGauge's idea to identify corner-cases in test data (*w.r.t.* training) among sets of neuron activations by singling out regions of neuron values that are never encountered in test data [24]. In our case, this naturally translates into the identification of hidden feature component intervals that are not exercised (enough) by the neuron values induced by the training dataset.

For each hidden feature component, our strategy for computing intervals using density estimates consists of a simple analysis of the associated density distribution. We define as interval boundaries the hidden feature values: (i) where the density crosses a given lower-bound threshold d_{min}; (ii) prominent local minima with density strictly greater than d_{min}. In Figure 3, the dashed vertical lines illustrate the result of such a discretisation strategy for the two considered hidden feature components.

Further, Simpler Discretisations   For the purposes of our initial approach, we also investigate two basic discretisation strategies that operate by partitioning each segment $\left[\min \left(\lambda_{i,j} \circ \hat{h}_i(X_{train}) \right), \max \left(\lambda_{i,j} \circ \hat{h}_i(X_{train}) \right) \right]$:

k-bins-uniform: The set of intervals partitions the segment into a given strictly positive number k of *bins*, all of the same width;

k-bins-quantile: The set of intervals is created as above, except that their respective width is calculated so that every interval holds a similar amount of individual projections from the training dataset $\lambda_{i,j} \circ \hat{h}_i(X_{train})$. For instance, the *4-bins-quantile* strategy proceeds by gathering all projected activations into 4 equally populated bins, and then computes 4 intervals based on the edges of each bin.

The resulting set of intervals is straightforwardly augmented with two left- and right-open intervals to obtain a partition of the hidden feature component line. These latter intervals are regions of neuron valuations that we assume correspond to corner cases for the DNN under investigation.

The combined set of extracted and discretised hidden features provides us with a *discrete* and *low-dimensional space* that we can use to reason about the neuron valuations induced by any given input x. This space will, in particular, allow us to associate each combination of hidden feature intervals with a measure of its occurrence within a given dataset (other that X_{train}), as well as reason about the dependencies between intervals of successive layers.

For the sake of conciseness, we will say that such an input (or set of inputs) *elicits* or *exercises* a given hidden feature interval $f_{i,j}^{\sharp k}$, noted $x \rightsquigarrow f_{i,j}^{\sharp k}$ (or $X \rightsquigarrow f_{i,j}^{\sharp k}$), when $\mathsf{Discr}_{i,j}^{\sharp} \circ \hat{h}_i(x) = f_{i,j}^{\sharp k}$. We generalise this notation to a given combination of intervals $F_i^{\sharp} = \left(f_{i,1}^{\sharp k_1}, \ldots, f_{i,|\Lambda_i|}^{\sharp k_{|\Lambda_i|}} \right) \in \mathbb{F}_i^{\sharp}$ for all hidden features extracted for a layer l_i as $x \rightsquigarrow F_i^{\sharp} \overset{\text{def}}{=} \bigwedge_{j \in \{1,\ldots,|\Lambda_i|\}} x \rightsquigarrow f_{i,j}^{\sharp k_j}$.

3.4 Construction of the Bayesian Network Abstraction

The abstraction that we construct primarily represents the *probabilistic distribution* of the set of hidden feature values induced by a test sample X. In other words, given an input $x \in X$, the abstraction allows us to estimate the probability that x induces a given combination of values for the hidden features that have been learned by the DNN.

Thanks to the layered and acyclic nature of the DNNs that we consider, we can directly characterise the *causal relationship* between the sets of neuron values in various layers *w.r.t.* a series of inputs as well. In other words, given an input $x \in X$, one can in principle estimate the conditional probability of each neuron value at layer l_i *w.r.t.* the probability of every combination of neuron values at layer l_{i-1}. By lifting the above relationship from individual neuron values to hidden feature intervals, we seek to capture *causal semantic relations* that link the features at each layer: in a layer l_i, and with an input x, the *probability* that a hidden feature valuation belongs to a given interval in the corresponding feature space is *dependent* on probabilities pertained to hidden feature intervals at layer l_{i-1}.

Let us assume that suitable feature-extraction and discretisation strategies have been employed to obtain a set of discretised feature components $\mathbb{F}_{\mathcal{N}}^{\sharp} = \left\{ \mathbb{F}_{1,1}^{\sharp}, \mathbb{F}_{1,2}^{\sharp}, \ldots, \mathbb{F}_{K,|\Lambda_K|}^{\sharp} \right\}$ for a network $\mathcal{N}$ with layers $(l_1, \ldots, l_K)$. We

first associate each hidden feature component $\mathbb{F}_{i,j}$ with a discrete (*i.e.,* categorical) random variable ranging over $\mathbb{F}_{i,j}^{\sharp}$. The causal relationship assumption further gives us the *structure* of a Discrete Bayesian Network, where each node is associated with a discretised hidden feature, and each edge denotes the dependency between a pair of hidden features in successive layers. Given this structure and a set of *inputs* $X \subset \mathbb{D}_X$, *Bayesian inference* on the sets of feature intervals $\left\{ \mathsf{Discr}_{i,j}^{\sharp} \circ \hat{h}_i(X) \,\middle|\, j \in \{1, \ldots, |\Lambda_i|\} \right\}$, for $1 \leq i \leq K$, allows us to estimate the probabilities mentioned above, and construct a BN $\mathcal{B}_{\mathcal{N},X} \stackrel{\text{def}}{=} (V_{\mathcal{N},X}, E_{\mathcal{N},X}, P_{\mathcal{N},X})$, where

$$V_{\mathcal{N},X} \stackrel{\text{def}}{=} \left\{ \emptyset f_{i,j}^{\sharp} \emptyset \,\middle|\, \mathbb{F}_{i,j}^{\sharp} \in \mathbb{F}_{\mathcal{N}}^{\sharp} \right\} \tag{6}$$

is the set of nodes, *i.e.,* one variable for each extracted feature component,

$$E_{\mathcal{N},X} \stackrel{\text{def}}{=} \bigcup_{1 < i \leq K} \left\{ \emptyset f_{i-1,1}^{\sharp} \emptyset, \ldots, \emptyset f_{i-1,|\Lambda_{i-1}|}^{\sharp} \emptyset \right\} \times \left\{ \emptyset f_{i,1}^{\sharp} \emptyset, \ldots, \emptyset f_{i,|\Lambda_i|}^{\sharp} \emptyset \right\} \tag{7}$$

is the set of edges connecting features in neighbouring layers, and

$$P_{\mathcal{N},X}\emptyset f_{i,j}^{\sharp}\emptyset \stackrel{\text{def}}{=} \begin{cases} \mathcal{P}_1\left(f_{1,j}^{\sharp k} \right) & \text{if } i = 1 \\ \mathcal{CP}_i\left(f_{i,j}^{\sharp k} \,\middle|\, F_{i-1}^{\sharp} \right) & \text{otherwise} \end{cases} \tag{8}$$

associates each node with either a marginal probability table (for hidden features of layer l_1) or a conditional probability table (for hidden or output layers). A conditional probability table $\mathcal{CP}_i\left(f_{i,j}^{\sharp k} \,\middle|\, F_{i-1}^{\sharp} \right)$ is defined for each feature interval $f_{i,j}^{\sharp k} \in \mathbb{F}_{i,j}^{\sharp}$ for layer l_i, *w.r.t.* each combination of feature intervals $F_{i-1}^{\sharp}$ for layer l_{i-1}, as

$$\mathcal{CP}_i\left(f_{i,j}^{\sharp k} \,\middle|\, F_{i-1}^{\sharp} \right) \stackrel{\text{def}}{=} \Pr\left(x \rightsquigarrow f_{i,j}^{\sharp k} \,\middle|\, x \rightsquigarrow F_{i-1}^{\sharp} \right). \tag{9}$$

Intuitively, $\mathcal{CP}_i\left(f_{i,j}^{\sharp k} \,\middle|\, F_{i-1}^{\sharp} \right)$ gives the probability that any input $x \in X$ chosen uniformly exhibits feature interval $f_{i,j}^{\sharp k}$, knowing that it exhibits intervals $F_{i-1}^{\sharp}$ for all extracted features at layer l_{i-1}.

One can compute *marginal probabilities* for hidden feature intervals that pertain to hidden and output layers based on the conditional probability tables as defined in Eq (9): the *unconditional probability* obtained for every interval $f_{i,j}^{\sharp k} \in \mathbb{F}_{i,j}^{\sharp}$ is then

$$\mathcal{P}_i\left(f_{i,j}^{\sharp k} \right) \stackrel{\text{def}}{=} \Pr\left(x \rightsquigarrow f_{i,j}^{\sharp k} \right). \tag{10}$$

Eq. (10) gives the probability that, given any input x chosen uniformly in X, $\lambda_{i,j} \circ \hat{h}_i(x)$ belongs to a given feature interval $f_{i,j}^{\sharp k}$, without any knowledge on the neuron values. Hidden features pertaining to the input later are not subject to any conditional dependence in the BN we construct. We therefore use $\mathcal{P}_1\left(f_{1,j}^{\sharp k} \right)$ in Eq. (8), which estimates the unconditional probabilities associated with hidden feature intervals of l_1.

Observe that, according to Eq. (7), the value of each feature extracted from a layer l_i is assumed to directly depend on the value of every feature for layer l_{i-1}; this assumption is further reflected in the construction of conditional probability tables in Eq. (8). Figure 2 contains an example BN, where the above assumption induces full connections between the nodes pertaining to successive layers. Figure 2 additionally gives two conditional probability tables for $\mathsf{Discr}^{\sharp}_{3,1} \circ h_3$ and $\mathsf{Discr}^{\sharp}_{3,2} \circ h_3$, respectively; it also shows two unconditional probability tables for $\mathsf{Discr}^{\sharp}_{1,1} \circ h_1$ and $\mathsf{Discr}^{\sharp}_{1,2} \circ h_1$, respectively. The tables are attached to their respective BN nodes.

3.5 Preserved Property

First of all, we show that the constructed BN is an abstraction of the NN. Given a finite set X of inputs, an abstraction constructs a set $X' \supseteq X$ that generalises X to more elements [12]. Given an input x and a BN $\mathcal{B}_{\mathcal{N},X}$, we are able to check the probability of x on $\mathcal{B}_{\mathcal{N},X}$, i.e., $\mathcal{B}_{\mathcal{N},X}(x)$. We say that x is included in $\mathcal{B}_{\mathcal{N},X}$ if $\mathcal{B}_{\mathcal{N},X}(x) > 0$. The following lemma suggests that every sample in X is included in $\mathcal{B}_{\mathcal{N},X}$:

Lemma 1. *All inputs x in the dataset X are included in the $\mathcal{B}_{\mathcal{N},X}$ with probability greater than 0.*

Let X' be the set of inputs that satisfy $\mathcal{B}_{\mathcal{N},X}(x) > 0$. This lemma suggests that $X' \supset X$. Therefore, $\mathcal{B}_{\mathcal{N},X}$ defines an abstraction of the dataset X. This abstraction also suggests that the abstraction assumption – *i.e.,* inputs that are outliers *w.r.t.* the abstraction are also outliers *w.r.t.* the original neural network – is reasonable because $X' \supset X$.

4 Exploiting the Bayesian Network

4.1 Probabilistic Inference

Given a BN $\mathcal{B}_{\mathcal{N},X}$, we can conduct probabilistic inference to understand how the original neural network processes inputs. By doing this, we gain a better understanding about the "black-box" neural network – and improve its interpretability/explainability – by conducting causal and evidential reasoning.

Causal Reasoning is to understand how the up-stream causes may affect the down-stream effects. Typically, it considers queries such as how the appearance of some input features may statistically affect how a given hidden feature or output label is exercised.

A BN is a joint probability $P(F^{\sharp}_1, \ldots, F^{\sharp}_K) = \prod_{i=1}^{K} \prod_{j=1}^{F^{\sharp}_i} P_{\mathcal{N},X}(\!|f^{\sharp}_{i,j}|\!)$—*cf.* Eq. (8). If we have evidence on the input feature by having, *e.g.,* $(\!|f^{\sharp}_{1,j_1}|\!) = v$ for some value v, then we will be able to reason about the maximum a posteriori

(MAP) estimation of the output labels, $MAP\left(\langle\!|f^{\sharp}_{K,j_2}|\!\rangle \,\middle|\, \langle\!|f^{\sharp}_{1,j_1}|\!\rangle = v\right)$, which represents the maximum probability the output label j_2 can have after considering not only the evidence $\langle\!|f^{\sharp}_{1,j_1}|\!\rangle = v$ but also all possible values of other nodes on the BN.

Evidential Reasoning is to understand how the availability of evidence on the down-stream effects may affect the up-stream causes. Typically, it considers queries such as how the appearance of some output label may affect the appearance probability of some input feature.

Considering a given label j_2 (that is, we have an evidence that $\langle\!|f^{\sharp}_{K,j_2}|\!\rangle = 1$ and $\forall j \neq j_2 \colon \langle\!|f^{\sharp}_{K,j}|\!\rangle = 0$), the probability of $\mathcal{P}_1\left(f^{\sharp k}_{1,j}\right)$ may change to reflect the availability of the evidence. It is possible that some features may be more important for some specific label.

4.2 Runtime Monitoring

We propose two ways of monitoring the behaviour of a DNN through its associated BN. The first one is to monitor the distribution shift by comparing a BN obtained from operational data and the original BN obtained using the training data. The second is to monitor the possibility of a given input to determine the risk of a safety violation by checking if it is an outlier.

Monitoring Distribution Shift Let $\mathcal{B}_{\mathcal{N},X} = (V_{\mathcal{N},X}, E_{\mathcal{N},X}, P_{\mathcal{N},X})$ be the BN we learned from the training dataset. It can be seen as a generative model for the dataset X. Covariate shift, an important kind of distribution shift, refers to the change in the distribution of the input variables present in the training and the test data. Now, assume that we have collected a set X' of operational data. We can use X' to construct another BN $\mathcal{B}_{\mathcal{N},X'} = (V_{\mathcal{N},X}, E_{\mathcal{N},X}, P_{\mathcal{N},X'})$ that has the same *structure* as $\mathcal{B}_{\mathcal{N},X}$ (*i.e.,* the vertices and the edges remain), but where the probability tables are re-generated according to X'.

By comparing the probabilities in $\mathcal{B}_{\mathcal{N},X}$ and $\mathcal{B}_{\mathcal{N},X'}$, we can assess whether some of the probabilistic causal relations are changed. Indeed, a change on probabilistic causal relation may imply the change on a conditional probability represented by the BN, on which the probabilistic causal relation is a factor.

Detecting Outliers Unlike the above methods which only consider the closeness of the new input with the training dataset in terms of geometric distance, we can assess whether any new input is an outlier by considering its probability on the BN: any input x' for which

$$\mathcal{B}_{\mathcal{N},X}(x') = 0 \tag{11}$$

will be identified as an outlier. As implied by Lemma 1, our abstraction is sound because all training instances are considered and none of them will lead to the case of Eq. (11).

5 Bayesian Network-based Feature Coverage

Given a trained DNN $\mathcal{N}$, the process described in Section 3.4 provides a method for constructing a BN $\mathcal{B}_{\mathcal{N},X}$, whose *structure* is derived using assumed dependencies between hidden features that are extracted and discretised by means of the techniques specified in Section 3.3. The probabilities in $\mathcal{B}_{\mathcal{N},X}$ are further inferred using the set of hidden feature valuations induced by X, and provide useful insights on whether X:

(i) induces a varied range of valuations for each extracted hidden feature; and
(ii) exercises the assumed relationships between hidden feature intervals.

Our BN abstraction can therefore be used to define some quality metric for the set X, in terms of the hidden features that are captured in layers of $\mathcal{N}$ that X does or does not exhibit. We elaborate on this idea in this Section, and develop new coverage metrics that can be used to assess the quality of test datasets. We show in the subsequent Section how these metrics can also be exploited to generate new inputs for the purpose of increasing the range of hidden feature valuations and causal relationships induced by such datasets.

5.1 Eliciting Semantic Assumptions Missed by Test Data

Observe that a marginal probability $\mathcal{P}_i\left(f_{i,j}^{\sharp k}\right) < \varepsilon$ computed based on a BN $\mathcal{B}_{\mathcal{N},X}$, where ε is a small-enough probability, expresses that $f_{i,j}^{\sharp k}$ is a hidden feature interval that is rarely elicited by any input in X. Similarly, a conditional probability $\mathcal{CP}_i\left(f_{i,j}^{\sharp k}\,\middle|\,F_{i-1}^{\sharp}\right) < \varepsilon$ expresses that an assumed causal dependency between intervals for all hidden features of layer l_{i-1} and a hidden feature interval $f_{i,j}^{\sharp k}$ for l_i is not exercised by X. Under the assumption that the hidden sub-space that underlies the structure of $\mathcal{B}_{\mathcal{N},X}$ reflects high-level (semantic) behaviours that $\mathcal{N}$ must capture, then both cases lead to clearly identified *semantic assumptions* that are not tested by any input from X, or not tested enough. We now turn these observation into more formal definitions of coverage metrics.

Hidden Feature Coverage The first BN-based coverage metric that we define disregards conditional probabilities, and concentrates on each individual hidden feature interval in isolation:

Definition 5.1 (BN-based Feature Coverage). Given a trained DNN $\mathcal{N}$, the *BN-based feature coverage* of a *non-empty* set of inputs $X \subset \mathbb{D}_X$ is obtained via the BN abstraction $\mathcal{B}_{\mathcal{N},X}$ as

$$\text{BFCov}\left(\mathcal{B}_{\mathcal{N},X}\right) \overset{\text{def}}{=} \frac{1}{\left|V_{\mathcal{N},X}\right|} \sum_{(\!\{f_{i,j}^{\sharp}\}\!) \in V_{\mathcal{N},X}} \frac{\left|\left\{f_{i,j}^{\sharp k} \in \mathbb{F}_{i,j}^{\sharp}\,\middle|\,\mathcal{P}_i\left(f_{i,j}^{\sharp k}\right) \geq \varepsilon\right\}\right|}{\left|\mathbb{F}_{i,j}^{\sharp}\right|}. \tag{12}$$

Informally, $\text{BFCov}\left(\mathcal{B}_{\mathcal{N},X}\right)$ ranges over $]0,1]$, and gives the proportion of features that are exercised (enough) by X—this measure cannot be null if ε is

sufficiently small, since the sum of all the entries of the probability tables is always one.

Example 5.1. Consider the BN given in Figure 2. For layer l_3, the following marginals can be computed based on the given two conditional probability tables: $\Pr(\lambda_{3,1} \circ h_3(x) < 3) \approx 0.725$, $\Pr(\lambda_{3,1} \circ h_3(x) \geq 3) \approx 0.275$, $\Pr(\lambda_{3,2} \circ h_3(x) < 3) \approx 0.4$, and $\Pr(\lambda_{3,2} \circ h_3(x) \geq 3) \approx 0.6$. Assuming a dataset X was used to infer the probabilities in the figure (and assuming similar non-negligible marginal probabilities for the nodes pertained to layers l_1 and l_2), then we obtain $\mathrm{BFCov}\,(\mathcal{B}_{\mathcal{N},X}) = 1$. This means that X adequately covers the full range of hidden features values learnt by the abstracted DNN.

Yet, the null conditional probabilities in the figure indicate that some combinations of hidden feature intervals between l_2 and l_3 are not exercised by X.

Hidden Feature-dependence Coverage   We further exploit the insights on causal relationships exercised by a dataset X that the conditional probabilities of a BN $\mathcal{B}_{\mathcal{N},X}$ provide to define the following coverage metric:

Definition 5.2 (BN-based Feature-dependence Coverage). Given a trained DNN $\mathcal{N}$, the *BN-based feature-dependence coverage* of a *non-empty* set of inputs $X \subset \mathbb{D}_X$ is obtained via the BN abstraction $\mathcal{B}_{\mathcal{N},X}$ as

$$
\mathrm{BFdCov}\,(\mathcal{B}_{\mathcal{N},X}) \overset{\mathrm{def}}{=} \frac{1}{\left| V_{\mathcal{N},X}^{+} \right|} \sum_{(f_{i,j}^{\sharp}) \in V_{\mathcal{N},X}^{+}} \frac{\left| \left\{ \begin{matrix} (f_{i,j}^{\sharp k}, F_{i-1}^{\sharp}) \in \\ \mathbb{F}_{i,j}^{\sharp} \times \mathbb{F}_{i-1}^{\sharp} \end{matrix} \middle| \begin{matrix} \mathcal{CP}_i\left(f_{i,j}^{\sharp k} \middle| F_{i-1}^{\sharp}\right) \geq \varepsilon \\ \vee \quad \mathcal{P}_i\left(f_{i,j}^{\sharp k}\right) < \varepsilon \end{matrix} \right\} \right|}{\left| \mathbb{F}_{i,j}^{\sharp} \times \mathbb{F}_{i-1}^{\sharp} \right|}
$$

(13)

where $V_{\mathcal{N},X}^{+} \overset{\mathrm{def}}{=} \left\{ (f_{i,j}^{\sharp}) \in V_{\mathcal{N},X} \middle| i > 1 \right\}$ are all nodes in $\mathcal{B}_{\mathcal{N},X}$ that represent hidden features extracted for hidden or output layers.

$\mathrm{BFdCov}\,(\mathcal{B}_{\mathcal{N},X})$ gives the proportion of assumed causal relationships between features of successive layers that are exercised (enough) by X. Observe that this metric only takes conditional probabilities into account; therefore, the first layer (*i.e.,* for $i = 1$) is taken out of the summation, and the normalisation factor is adjusted accordingly. Further, the set above the innermost fraction includes the conditional probability entries pertaining to hidden feature intervals that are *not* elicited by X (*i.e.,* for which $\mathcal{P}_i\left(f_{i,j}^{\sharp k}\right) < \varepsilon$). We eliminate the impact of feature coverage as measured in Def. 5.1 on feature-dependence coverage in this way, *i.e.,* so that $\mathrm{BFdCov}\,(\mathcal{B}_{\mathcal{N},X})$ only captures the causal relationships between hidden features that are not exercised by X, whatever the actual coverage of hidden features that X induces.

At last, we straightforwardly combine the two above metrics to provide a consistent coverage measure that is based on every probability entry of the BN:

$$
\mathrm{BFxCov}\,(\mathcal{B}_{\mathcal{N},X}) \overset{\mathrm{def}}{=} \mathrm{BFCov}\,(\mathcal{B}_{\mathcal{N},X}) \times \mathrm{BFdCov}\,(\mathcal{B}_{\mathcal{N},X}).
$$

(14)

Example 5.2. Continuing Example 5.1 (and, again, further assuming non-negligible entries in the conditional probability tables that are not shown in Figure 2), we obtain

$$\text{BFdCov}\left(\mathcal{B}_{\mathcal{N},X}\right) = \frac{1}{4} \times \left(2 \times \frac{8}{8} + 2 \times \frac{7}{8}\right) = \frac{15}{16} \approx 94\%.$$

Coverage Criteria We trivially derive the following test criteria from the coverage metrics given in Defs. 5.1 and 5.2:

Definition 5.3 (BN-based Feature Coverage Criterion). A non-empty set of inputs $X \subset \mathbb{D}_X$ satisfies the *BN-based feature coverage criterion* that is obtained via the BN abstraction $\mathcal{B}_{\mathcal{N},X}$ iff $\text{BFCov}\left(\mathcal{B}_{\mathcal{N},X}\right) = 1$.

Definition 5.4 (BN-based Feature-dependence Coverage Criterion). A non-empty set of inputs $X \subset \mathbb{D}_X$ satisfies the *BN-based feature-dependence coverage criterion* that is obtained via the BN abstraction $\mathcal{B}_{\mathcal{N},X}$ iff $\text{BFdCov}\left(\mathcal{B}_{\mathcal{N},X}\right) = 1$.

5.2 Advantages of the Abstraction and Coverage Criteria

In the following, we discuss a few advantages of our BN-based approach to modelling DNN behaviours, and the adjoined coverage criteria.

From Structural Coverage to Semantics Coverage Our BN-based abstraction allows us to formulate semantics-based coverage criteria: we are now able to work directly with features instead of neurons. As argued in the introduction, in machine learning, features are regarded as the basic elements and should be dealt with directly.

Working with Semantics with Minimal Adaptation Definitions 5.3 and 5.4 resemble structural versions of coverage criteria, that are directly based on neuron activations [26, 24, 32]. Actually, $\text{BFCov}\left(\mathcal{B}_{\mathcal{N},X}\right)$ can be seen as the BN counterpart of the neuron coverage [26] and $\text{BFdCov}\left(\mathcal{B}_{\mathcal{N},X}\right)$ can be seen as the BN counterpart of the MC/DC [32]. Therefore, our BN representation of hidden space coverage provides us with a means to enhance existing structural approaches for assessing the quality of trained DNNs with semantics, without requiring drastic changes to the basic principles. We will demonstrate this aspect further in the next Section, by showing how the core test case generation algorithm of DeepConcolic can be adapted to the above criteria.

Evaluating Suitability of Feature Extraction Techniques on a Dataset Our framework is general and can work with any feature extraction technique. Therefore, it can potentially be applied to evaluate and compare the quality of applying feature extraction techniques on some given dataset—it is known that there does not exist a single feature extraction technique that works best on all datasets. Certainly, to have a fair evaluation and comparison, the dataset needs to be of "golden standard".

Improved Scalability The size—measured by the number of nodes—of the BN is adjustable, and usually significantly smaller than the original neural network. However, as hinted in Eq. (8), the size of the underlying conditional/joint probability tables at a given layer l_i grows exponentially in the number of nodes used to represent layer l_{i-1} (*i.e.,* its number of hidden features) and in the size of their respective partitions.

Still, if the number of extracted hidden features and their respective partitioning is kept small enough, then the scalability of any analysis method based on our BN representation can be significantly better than an analysis method that is directly based on the neural network's structure.

Connection to Boxing-Clever, and Beyond The production of the marginal probability tables of the input layer can partition the input space into a set of hyper-rectangles. Each hyper-rectangle is associated with a set of constraints (or half-planes), such that each constraint comes from one hidden feature of the input layer. Therefore, by working with the marginal probability tables, it is easy to work out some functionalities that are enabled in Boxing-Clever [1], *e.g.,* the computation of **empty hyper-rectangles**. Note that in our case, the hyper-rectangles are based on hidden features instead of neurons.

Beyond the input layer, every 0-valued entry of a conditional probability table (of some hidden layer) represents a missing training example in satisfying the causal relation represented by the 0-valued entry. By back-propagation of the 0-value to the input layer, it is possible to identify an **empty region** – not necessarily an empty hyper-rectangle – of the input space such that any input in the region satisfies the causal relation represented by the 0-valued entry.

Detection of Unintended Memorisation and its Implication to Safety Assurance Intuitively, if the test dataset X is deemed representative of realistic inputs, then low-probability entries in the Bayesian Network $\mathcal{B}_{\mathcal{N},X}$ may reveal out-of-distribution training data samples. Carlini et al. [4] have suggested that a neural network may unintendedly memorise these out-of-distribution data samples, leading to a negative effect on its generalisation ability [7]. We have discussed in [40, 41] the relation between generalisation ability with the failure rate of neural networks in operation, under a safety assurance framework.

6 Concolic Approach for Test Case Generation

We have extended the DeepConcolic approach for testing DNNs that was pioneered by Sun et al. [32] with our new test criteria in order to operationally assess the utility of the above constructions and definitions. Our goal in doing so is to check whether our new definitions for hidden feature coverage can be used to guide the generation of test cases.

6.1 DeepConcolic in a Nutshell

The DeepConcolic approach for test case generation consists of alternating *concrete* executions and *symbolic* analyses of the DNN under test. The concrete executions actually evaluate the test input using the trained DNN, and report on the level of coverage achieved. The symbolic analyses, on the other hand, aim at synthesising new test inputs based on some *test target* that is chosen so as to increase the considered coverage metric. This approach is provably more efficient at constructing sets of test inputs that achieve very good coverage than a basic random generation of test inputs.

Test Case Generation via Symbolic Analysis   The symbolic analysis is performed on a symbolic model that constrains every neuron activation up to a layer l_i according to the functional semantics of each respective layer, with the addition of neuron valuation constraints that encode a given test target (such as a change in the sign of a particular neuron output, for instance). The goal of the analysis is then to solve the problem of finding a value for every input neuron, such that some distance metric *w.r.t.* a legitimate test input is minimised, while satisfying all the above constraints. In the case of input images, typical distance metrics for assessing the similarity between inputs include the Chebyshev distance L_∞, which uses the maximum absolute difference between any one pixel or colour component between two images as a measure of their distance, or the L_0 "norm", that simply counts the number of pixels or colour components that differ. The set of values obtained for every input neuron consists of a new test case that is in principle guaranteed to fulfil the test target; this input is added to the set of generated test cases if some *oracle* mechanism determines that it is "close enough" to a reference test input. The latter requirement serves as a "plausibility" filter to only retain realistic inputs.

Linear Programme and Distance Metric   The problem that is solved at each symbolic analysis step of DeepConcolic is a Linear Programme (LP), and in our case, the symbolic analysis consists in solving an LP as well. Such a problem is built as a set of *constraints Constr* $(\cdot)$ that involve continuous (free) *variables*, associated with an *optimisation objective*. A solution to the problem consists of an assignment for every one of these variables that meets both the set of constraints and the optimisation objective.

The specification of an LP problem requires that every constraint and objective function involved be *linear inequalities* or *linear arithmetic expressions*. This requirement constrains the distance metric that we can use to compare input images, which must therefore be linear as well. Among the metrics listed above, only the Chebyshev distance L_∞ satisfies this requirement. We will therefore assume this distance metric below, and will denote with $\|x - y\|_\infty$ the distance between two vectors of same dimension x and y.

Convergence & Approximations   In some cases, reliance on a linear encoding of functional semantics may induce losses in analysis precision. This

translates in either:

over-approximation: incorrect solutions, *i.e.,* the LP-problem over-approximates the actual behaviours of the DNN on the given input;

under-approximation: infeasible problems derived from legitimate behaviours, *i.e.,* the LP-problem under-approximates the actual functional behaviours of the DNN on the given input.

In DeepConcolic, such an imprecision typically arises in layers whose functional semantics cannot be precisely encoded using linear arithmetic and inequalities. Such is the case for max-pooling layers for instance, that require *ad hoc* treatments for its encoding as a set of linear constraints.

One can easily check whether the input that is extracted from a solution to the LP problem exhibits the intended behaviours so as to discard inputs that do not meet the desired test target; keeping such inputs may also help better exploring the input-space. Note, however, that both under- and over-approximations may impair the convergence of the overall generation algorithm. We elaborate on this particular aspect for our adaptation in the next Section, and investigate it further experimentally in Section 7.2.

6.2 Concolic Test Generation with BN-based Criteria

In our case, a test target, that we will denote $\mathit{Target}\left(f_{i,j}^{\sharp k}\right)$, consists of a (set of) hidden feature interval(s) that should be elicited by the test input to be generated; such is the case for the criterion given in Def. 5.3 above for instance, for covering a rare hidden feature interval.

The overall procedure is parameterised with the structure of the Bayesian Network based on which our coverage metrics are defined: *i.e.,* we assume here that suitable feature extraction and discretisation have been applied on a training sample X_{train}. The computation starts by randomly sampling an initial set of test inputs X_0 that is correctly classified by $\mathcal{N}$, and initialising the probability tables in the Bayesian Network to produce $\mathcal{B}_{\mathcal{N},X_0}$. Let $X_0^{ok} = X_0$ be the initial set of candidates from which new tests can be derived. Then, each iteration i of the test case generation process proceeds according to the following operations:

1. Identify a test target $t = \mathit{Target}\left(\cdot\right)$ that is not yet met by the current set of input test cases X_i. This step is performed by means of an analysis of the marginal or conditional probability tables in $\mathcal{B}_{\mathcal{N},X_i}$;
2. Select a test input $x \in X_i^{ok}$ according to some heuristics, such as some closeness to the identified target t;
3. Construct an LP problem based on t: this problem comprises a set of constraints *Constraints*, and an optimisation objective that seeks to minimise the distance between activations of input neurons and x. This problem is formulated as:

$$\text{Minimise: } \|(n_{1,1}, \ldots, n_{1,|l_1|}) - (x_{1,1}, \ldots, x_{1,|l_1|})\|_\infty \tag{15}$$
$$\text{Subject to: } \mathit{Constraints}$$

where $n_{1,1}, \ldots, n_{1,|l_1|}$ is the set of all input neurons in $\mathcal{N}$;

4. Solve the LP problem, and extract the newly generated test input x' from values of input neurons: $x' = (n_{1,1}, \ldots, n_{1,|l_1|})$;

5. Keep the generated input x' if it passes the oracle *and, optionally, if it actually improves on x w.r.t. the target t*: in such a case, let $X_{i+1} = X_i \cup \{x'\}$; otherwise, let $X_{i+1} = X_i$ and $X_{i+1}^{ok} = X_i^{ok}$, and continue from step 1;

6. Let $X_{i+1}^{ok} = X_i^{ok} \cup \{x'\}$ if $f_{\mathcal{N}}(x') = f_{\mathcal{N}}(x)$, X_i^{ok} otherwise. The latter case indicates that $\mathcal{N}$ does not output the same classification label for x' than for x, which means that x' is considered *adversarial* for $\mathcal{N}$, as x' is both deemed close enough to x from which it is derived, and it is not assigned the same label as x by $\mathcal{N}$;

7. Update probabilities in $\mathcal{B}_{\mathcal{N}, X_i}$ to account for the new test x' and construct $\mathcal{B}_{\mathcal{N}, X_{i+1}}$; the test case generation continues if the test criterion obtained via this new BN is not satisfied.

The three steps from the algorithm above that need to be specialised according to the sought-after BN-based test criteria consist of:

(i) the identification of a test target by analysing $\mathcal{B}_{\mathcal{N}, X_i}$;

(ii) the selection of a test input;

(iii) the construction of the set of constraints *Constraints*.

Test oracle aside, the selection criterion in step 5 embodies a practical measure that we employ to counter a consequence of the loss of precision that is induced by dimensionality reduction. We discuss this aspect further in Section 6.3 below. We first elaborate on step 3, which involves an encoding of the functional behaviours of the DNN up to some layer of interest, that is indicated by the test target.

6.2.1 Encoding the Neural Network

Given a test target that aims at obtaining a test input that exhibits a given hidden feature interval at layer l_k, we first construct a symbolic encoding of the neural network up to l_k. We reuse the constraints suggested by Sun et al. [33] to construct this encoding of layer behaviours. Precisely, for each neuron $n_{i,j}$ for $i > 1$, and assuming ReLU activation functions, we have two variables $\hat{n}_{i,j}$ and $n_{i,j}$ and constraints that can directly be derived from Eqs. (1) and (4) in Section 2:

$$\hat{n}_{i,j} = W_{i,j} \left(n_{i-1,1}, \ldots, n_{i-1,|l_{i-1}|} \right) + b_i,$$
$$n_{i,j} \geq \hat{n}_{i,j} \qquad \text{and} \qquad n_{i,j} \geq 0$$

where $W_{i,j}$ and b_i are weight and bias parameters in layer i, respectively.

The case of layers that induce non-linear functional behaviours requires an alternative approach, since they cannot be directly encoded in the same way as in Eq. (6.2.1). To circumvent this issue, our approach consists in constraining

the neurons of these layers in such a way that they exhibit the same behaviour as that induced by the candidate input x. Considering for instance a max-pooling layer l_i (for which we gave the functional behaviour in Eq. (2) in Section 2), the idea is first to capture its selection behaviour using the following binary matrix $S^x \in \{0,1\}^{|l_i| \times |l_{i-1}|}$, whose coefficients are computed as

$$S^x_{j,k} = (1 \text{ if } \hat{n}_{i,j} = n_{i-1,k}, 0 \text{ otherwise})$$

so that a 1 at row j and column k indicates that the output of neuron $n_{i-1,k}$ as induced by x was selected by the layer as the value for $\hat{n}_{i,j}$, $i.e.$, $n_{i-1,k} = \max s_j \left(n_{i-1,1}, \ldots, n_{i-1,|l_{i-1}|} \right)$ as per Eq. (2). Note that multiple neurons of layer l_{i-1} selected by a function s_j may be maxima, in which case there are more than one 1's in the corresponding row of S^x. Then, obtaining a new input that induces the same selection behaviour at layer l_i boils down to enforcing either one of the following linear constraints:

$$(\hat{n}_{i,j} = n_{i-1,k}) \text{ if } S^x_{j,k} = 1, (\hat{n}_{i,j} > n_{i-1,k}) \text{ otherwise.}$$

For an input neuron $n_{1,j}$, we have a variable $n_{1,j}$ and a constraint on its possible values based on the range of each input component ($e.g.$, pixel)

$$n_{1,j} \in [a, b].$$

We let $Constr_{1,\ldots,i} (\mathcal{N})$ be the set of constraints that we obtain as described above for the neural network $\mathcal{N}$ up to layer l_i.

6.2.2 Targeting Feature Intervals

Further, a test target $Target \left(f^{\sharp k}_{i,j} \right)$, that aims at eliciting hidden feature interval $f^{\sharp k}_{i,j} \in \mathbb{F}^{\sharp}_{i,j}$ in layer l_i, directly translates into a pair of inequalities on neuron values for layer l_i as follows:

$$Constr \left(f^{\sharp k}_{i,j} \right) \stackrel{\text{def}}{=} \left(f^{[k}_{i,j} \leq \lambda_{i,j} \left(\hat{n}_{i,1}, \ldots, \hat{n}_{i,|l_i|} \right) < f^{k[}_{i,j} \right) \tag{16}$$

Observe that, if the feature mapping $\lambda_{i,j}$ is linear, then the inequalities above are linear as well. We straightforwardly extend the above construction to sets of target feature intervals for a layer l_i: $Target \left(F^{\sharp}_i \right)$ translates into $Constr \left(F^{\sharp}_i \right)$, that is the union ($i.e.$, the conjunction) of all inequalities built as per Eq. (16) for each interval in $F^{\sharp}_i$.

6.2.3 Fulfilling BN-based Test Criteria

When the goal of the symbolic analysis is to increase BFCov $(\mathcal{B}_{\mathcal{N},X})$ coverage, the identification of a test target consists in finding an interval of hidden feature values that has a small probability of occurring according to $\mathcal{B}_{\mathcal{N},X}$: this search

boils down to identifying a target interval $f_{i,j}^{\sharp k}$ such that $\mathcal{P}_i\left(f_{i,j}^{\sharp k}\right) < \varepsilon$, and building a target $t = \mathit{Target}\left(f_{i,j}^{\sharp k}\right)$.

Finding a candidate test input that has the best chance of leading to an interesting solution for the LP problem is a hard problem. Even more so when considering the potential losses in precision incurred by the underlying symbolic analyses. We therefore make the simple assumption that an input x whose feature value $\lambda_{i,j} \circ \hat{h}_i(x)$ is close to any of the target interval boundaries is a "good-enough bet". This assumption gives a simple heuristic for finding the candidate test input x, which consists in searching for a test $x \in X$, that minimises either $|\lambda_{i,j} \circ \hat{h}_i(x) - f_{i,j}^{[k}|$ or $|\lambda_{i,j} \circ \hat{h}_i(x) - f_{i,j}^{k[}|$.

Then, the set of constraints *Constraints* for specifying the LP problem for targeting t consists of:

$$\mathit{Constraints} = \mathit{Constr}_{1,\dots,i}\left(\mathcal{N}\right) \cup \mathit{Constr}\left(f_{i,j}^{\sharp k}\right) \tag{17}$$

The case of BFdCov $\left(\mathcal{B}_{\mathcal{N},X}\right)$ is similar, as the test target identification consists in finding an entry $\mathcal{CP}_i\left(f_{i,j}^{\sharp k} \,\middle|\, F_{i-1}^{\sharp}\right) < \varepsilon$ in any conditional probability table $P_{\mathcal{N},X}(f_{i,j}^{\sharp})$. One straightforwardly obtains $t = \mathit{Target}\left(\left\{f_{i,j}^{\sharp k}\right\} \cup F_{i-1}^{\sharp}\right)$, that aims at covering both the interval $f_{i,j}^{\sharp k}$ and a condition on hidden features of layer l_{i-1} for which the probability of finding an input in X that exhibits $f_{i,j}^{\sharp k}$ is very small. The selection of a test input can also be done according to a heuristic search like the one above, and the resulting set of constraints is:

$$\mathit{Constraints} = \mathit{Constr}_{1,\dots,i}\left(\mathcal{N}\right) \cup \mathit{Constr}\left(F_{i-1}^{\sharp}\right) \cup \mathit{Constr}\left(f_{i,j}^{\sharp k}\right) \tag{18}$$

6.3   Helping Convergence

In the description of the algorithm above, we have hinted at a practical measure that we use to counter a consequence of the loss of precision that is induced by dimensionality reduction. Indeed, although the latter comes in addition to the imprecision that is notably due to linearly encoding non-linear behaviours such as max-pooling layers—discussed in Section 6.1—, our preliminary empirical evaluations have shown us that the aforementioned imprecision has an adverse effect on the ability of our adapted algorithm to fulfil our BN-based criteria.

For instance, given a test target $t = \mathit{Target}\left(f_{i,j}^{\sharp k}\right)$ and a candidate test input x, the solution x' of the resulting LP program often leads to a hidden feature valuation $\lambda_{i,j} \circ \hat{h}_i(x')$ that does not belong to $f_{i,j}^{\sharp k}$: in other words, the actual behaviour at l_i subject to input x' does not satisfy Eq. (16) due to the approximations induced by the feature mapping $\lambda_{i,j}$. Even worse, x' may actually be "further" on the $\mathbb{F}_{i,j}$ line from the target than x: *i.e.,* $|t - \hat{h}_i(x)| <$ $|t - \hat{h}_i(x)|$, where $|t - y| \stackrel{\mathrm{def}}{=} \min\left(|\lambda_{i,j}(y) - f_{i,j}^{[k}|, |\lambda_{i,j}(y) - f_{i,j}^{k[}|\right)$ measures the

24

distance between t and the hidden feature value induced by neuron values y for layer l_i.

On a purely theoretical level, we acknowledge that this convergence issue is an intrinsic limitation of our technique, as it may be possible that none of the constructed LP problems ever have any solution that improves the desired coverage metric. Still, we can identify ways to counter this issue in practice. First, we can observe that rejecting any newly generated input on the grounds that it does not fulfil the test target quickly leads to the inability to find new candidate test inputs for any unmet test target. Instead, considering such inputs as legitimate new test cases even though they do not improve coverage helps populating the set of potential candidate tests inputs: they indeed help further explorations of the input space. Further, keeping only a new input x' if it improves over its original test x *w.r.t.* the target t may naturally lead to the construction of a set of test cases that is "closer" to satisfying the criterion (although this is not a requirement for actually helping convergence). We validate these claims further with dedicated experiments in Section 7.1.

At last, we posit that further constraining the set of admissible solutions of the LP programs indirectly gears the set of generated inputs towards cases that meet the desired targets. In particular, we assume that constraining non-targeted hidden features so that they replicate the respective values induced by x reduces the over-approximations, *i.e.,* increases the precision of the symbolic analysis. We achieve this by introducing the additional set of linear constraints

$$Replic_x \left(f_{i,j}^{\sharp k} \right) \stackrel{\text{def}}{=} \bigcup_{j' \in \{1, \ldots, |\Lambda_i|\}, j' \neq j} \left(\lambda_{i,j'} \left(\hat{n}_{i,1}, \ldots, \hat{n}_{i,|l_i|} \right) = \lambda_{i,j'} \circ \hat{h}_i(x) \right) \quad (19)$$

into the set *Constraints* built for targeting $Target \left(f_{i,j}^{\sharp k} \right)$ from an input x (in the case of a BN-based feature criterion). We empirically support this assumption using a dedicated set of experiments in the next Section.

7 Implementation and Evaluations

We have augmented the DeepConcolic tool[3] in order to experimentally validate the practicality and efficiency of our new coverage metrics, as well as our adaptation of concolic test case generation. With the following experiments, we want to assess whether our BN-based approach meets the needs to define semantics-based coverage criteria, that go beyond the low-level structure of neural networks.

We have implemented multiple strategies for linear dimensionality reduction and discretisation of each feature component. We have notably included PCA (with or without pre-scaling of activation values) and ICA (*cf.* Section 2), as well as several strategies for computing discrete partitions (*cf.* Section 3.3).

The wide range of combinations that all these strategies offer, even more when a distinct choice can be made regarding each individual layer, makes the

---

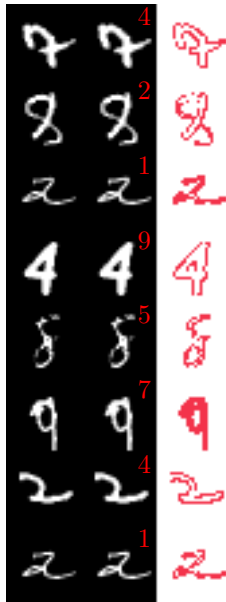[3]Available at https://github.com/TrustAI/DeepConcolic/.

Figure 5: Adversarial examples found by achieving a BN-based feature coverage criterion for all dense layers of $\mathcal{N}_{\mathsf{ms}}$. Each row includes, from left to right, the original test image, the adversarial example (with its new classification label in overlay), and an image that emphasises every pixel that differs.

complete study of this choice a challenging task. As in the first step we mainly want to assess the practicality of our BN-based approach for defining coverage criteria, for now we will report on some results for broad families of combinations of strategies. All experiments were carried out on the DNNs $\mathcal{N}_{\mathsf{ms}}$ and $\mathcal{N}_{\mathsf{mx}}$ described in Tables 1 and 2, that target the classification of the handwritten digits from the MNIST dataset.

7.1 Concolic Testing for BN-based Criteria

We have run the extended concolic testing tool on $\mathcal{N}_{\mathsf{ms}}$ and $\mathcal{N}_{\mathsf{mx}}$ by focusing on their dense layers (named `dense` and `dense_1` in both cases), with varying numbers of extracted feature components for each one of these two layers, using various algorithms for feature extraction, and/or changing the discretisation strategies. Each run comprised at most 100 iterations, and was initialised with uniformly drawn test sets X_0 of 10, 100, or 1000 correctly classified inputs.

We give in Figure 5 some adversarial examples that were found during the experiments on $\mathcal{N}_{\mathsf{ms}}$ and when targeting the BN-based feature coverage criterion. For comparison purposes, we also show in Figure 6 an adversarial example found by DeepConcolic to achieve a structural criterion (*i.e.,* expressed on individual neuron activations instead of on hidden features). We observe that,
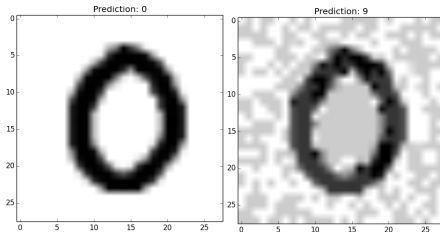
26

Figure 6: Adversarial example found by DeepConcolic while seeking a structural coverage criterion; this is borrowed from Sun et al. [33].

even though the distance metric used in all three cases was the same (L_∞), the example found through structural coverage appears to involve small-grain perturbations of pixels. On the contrary, our adversarial examples, and in general every generated test, exhibit differences that seem to precisely indicate where the input images need to be perturbed to alter the induced value for the respective targeted hidden feature.

We note that most adversarial examples were found during the first 30 to 40 iterations of the test case generation algorithm. Further, most adversarial examples were found when using *extended* discretisation strategies (*cf.* Section 3.3), where two additional left- and right- open intervals that do not contain any value that is induced by the training dataset are used for partitioning hidden feature components. This suggests that attempting to identify intervals in a hidden feature component $\mathbb{F}_{i,j}$ that do not exhibit any feature value (*i.e.,* finding intervals whose intersection with the set $\lambda_{i,j} \circ \hat{h}_i(X_{train})$ is empty, where X_{train} is the training dataset), seems a suitable strategy to find corner-cases for DNNs.

Let us now turn to a discussion on the overall behaviours of the concolic testing approach delineated in Section 6 in regard to the BN-based coverage metrics defined in Section 5. In particular, we want to assess the impact of the size of the initial test set (X_0) on the ability of DeepConcolic to generate new sets of inputs that achieve high coverage.

We summarise the runs for BN-based feature coverage in Figure 7. We first observe that, overall, between 10% to 60% of iterations produce new test inputs, and that the increase in the size of X_i is not correlated with the corresponding feature coverage. Instead, the coverage tends to increase in discrete steps, and often reaches plateaus for significant numbers of iterations. The former observation is a direct consequence of the discrete nature of the structure—the BN—based on which we defined the coverage metrics. This behaviour also gives a great illustration of our concerns for convergence raised in Section 6.3. We also note that neither PCA (in blue) or ICA (in red) seem to have any impact on the number of generated tests or the achieved coverage.

Further, the bottom row in Figure 7 indicates that 100 iterations already leads to increases in achieved coverage when the initial size of the test set X_0 is sufficiently large. This confirms that the size of the test set helps our heuristics
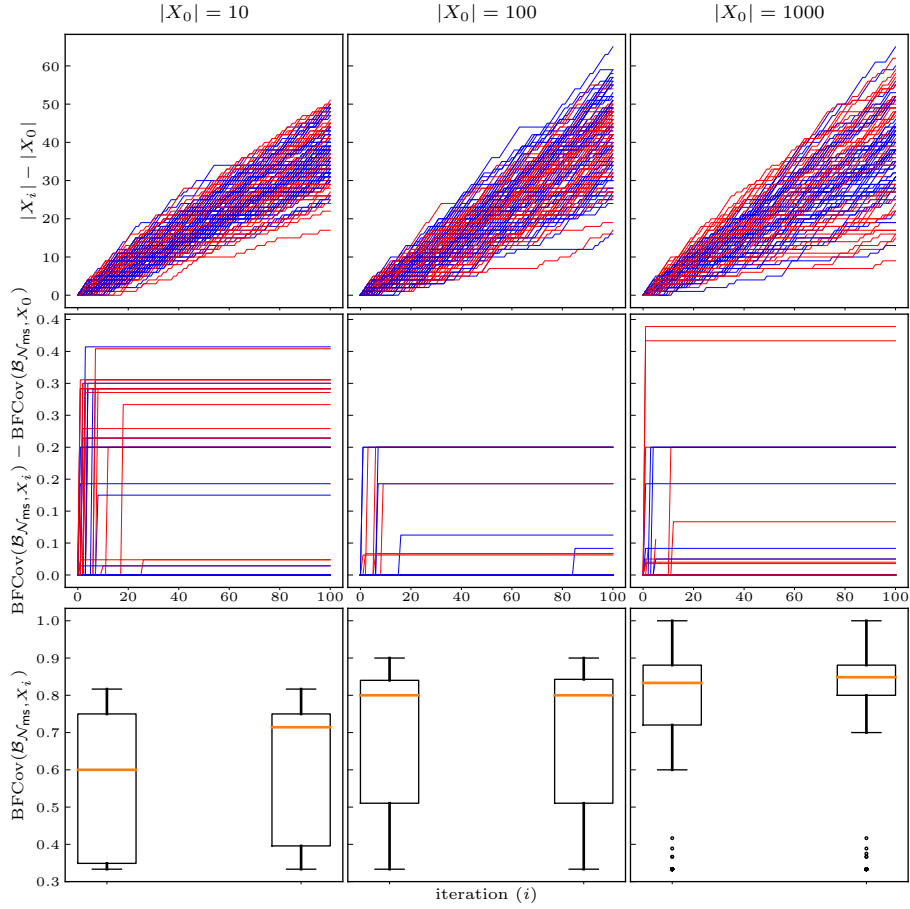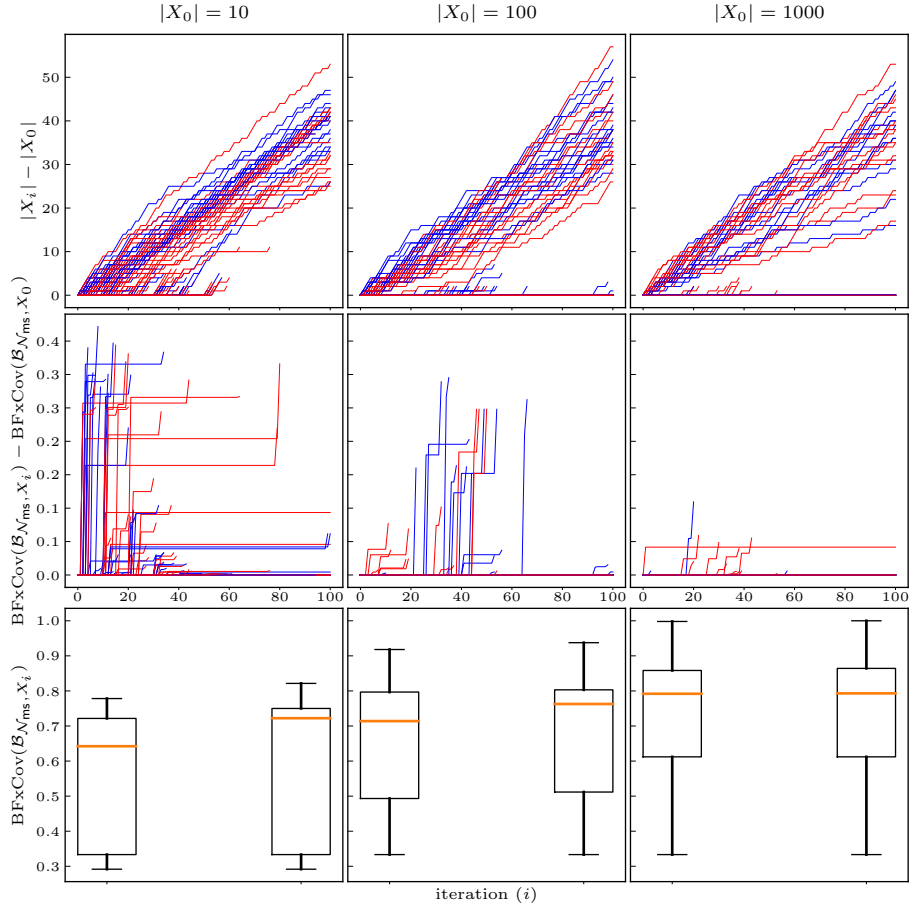
Figure 7: Representation of 360 traces of up to 100 iterations of test case generation by DeepConcolic targeting BN-based feature coverage, for various sizes of initial test sets $|X_0| \in \{10, 100, 1000\}$. Red and blue lines respectively indicate runs with ICA and PCA-based feature extractions. Each box in the bottom row shows the overall distribution of initial BFCov $(\mathcal{B}_{\mathcal{N}_{\mathsf{ms}}, X_0})$ and the respective final coverage BFCov $(\mathcal{B}_{\mathcal{N}_{\mathsf{ms}}, X_{100}})$.

Figure 8: Representation of 360 traces of up to 100 iterations of test case generation by DeepConcolic targeting BN-based feature-dependence coverage, for various sizes of initial test sets $|X_0| \in \{10, 100, 1000\}$. Red and blue lines respectively indicate runs with ICA and PCA-based feature extractions. Each box in the bottom row shows the overall distribution of initial BFxCov $(\mathcal{B}_{\mathcal{N}_{\mathsf{ms}}, X_0})$ and the respective final coverage BFxCov $(\mathcal{B}_{\mathcal{N}_{\mathsf{ms}}, X_{100}})$.

in identifying good candidate test inputs for a given test target.

We now turn to our results for BN-based feature-dependence coverage, summarised in Figure 8; note that the coverage measure that is plotted is the one defined in Eq. (14), which is the combined version of pure feature coverage and feature-dependence. We first observe that many runs terminate before reaching 100 iterations. This can be tracked down to two root causes: (i) the feature-dependence coverage reaches 100% (while full feature coverage is not achieved), which basically means that every conditional probability table entry is greater than ε whereas some entries of marginal probability tables are still negligible; (ii) no new pair (test target, candidate test input) can be identified (*i.e.,* every such pair has already been attempted). Still, we can note an overall increase in coverage for $|X_0| = 10$. This suggests that a dedicated heuristics is required to achieve both feature and feature-dependence coverage in combination.

7.2 Convergence & Precision Aspects

In Section 6.3, we posited that constraining the LP program in such as way that hidden feature components that are not related to the target replicate the values induced by the source input x, alleviates the convergence issues. We have again used $\mathcal{N}_{\mathsf{ms}}$ to assess this assumption. We have conducted two sets of experiments, that each consisted in a set of 40 runs with each parameter uniformly drawn from a set of pre-defined values for the number of extracted features per layer, discretisation strategy, number of bins, etc. Each set of runs differs in that for a single one of them the counter measure detailed in Eq. (19) was employed.

We plot in Figure 9 the resulting distributions of: distance from generated inputs to target (d — on the respective components line), along with the similar histograms showing the difference between the distance of a source test input x and a newly generated one x': $\delta > 0$ indicates that the generated input is closer to the target hidden feature interval than its source test. We can mainly observe that using the counter measure has no noticeable impact on our measure of progress (δ), yet it seems to gear generated inputs closer to their respective targets overall. We made similar observations when PCA was used for feature extraction, and when performing the same experiments with other DNNs.

We could note, however, that the above impact seems to be barely noticeable in the case of networks with max-pooling layers. We illustrate this by showing in Figure 10 the histograms obtained from an identical set of experiments for $\mathcal{N}_{\mathsf{mx}}$. This suggests that the benefits of the counter measure is attenuated by the loss in precision that is incurred by our symbolic encoding of such intrinsically non-linear layers.

8 Related Work

The investigation of techniques to analyse the weaknesses of deep learning has been a very popular research domain in the past few years. In the following,
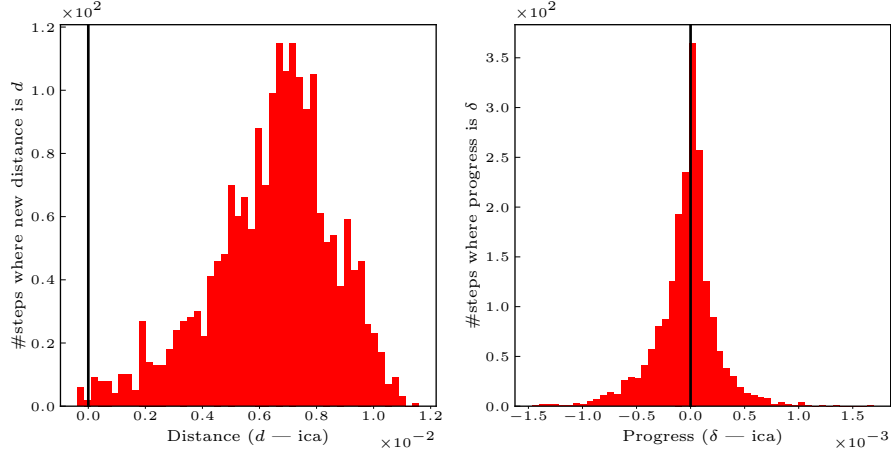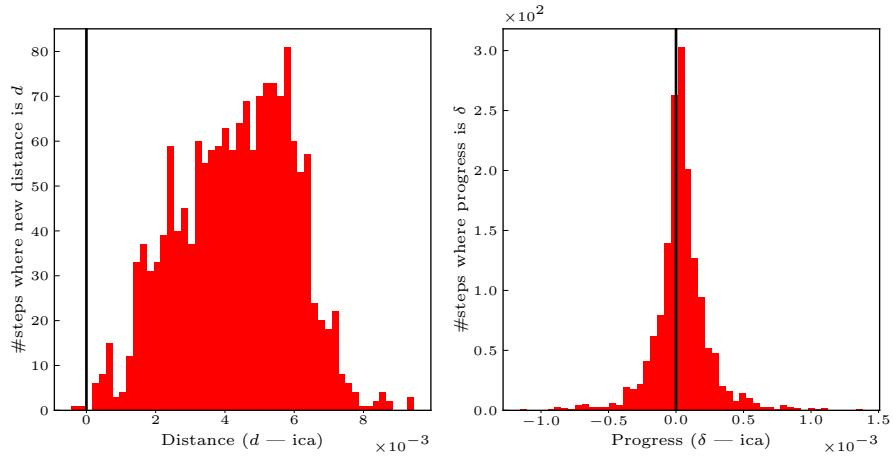
(a) without replication of values for non-targeted hidden feature components
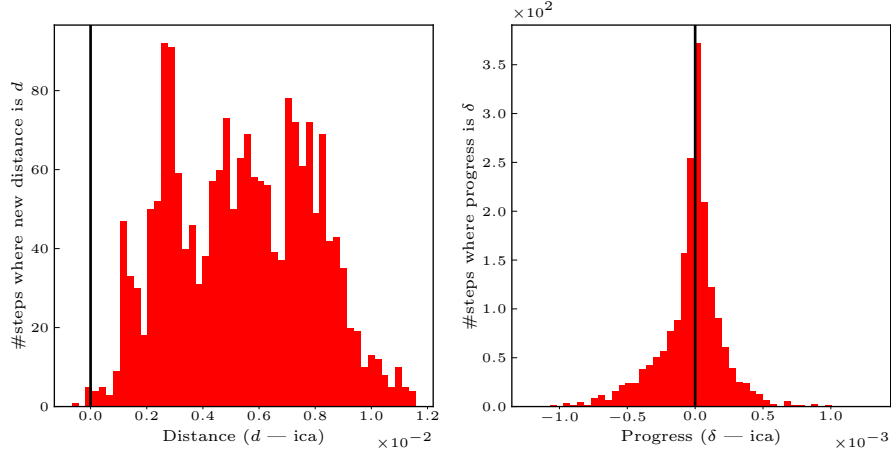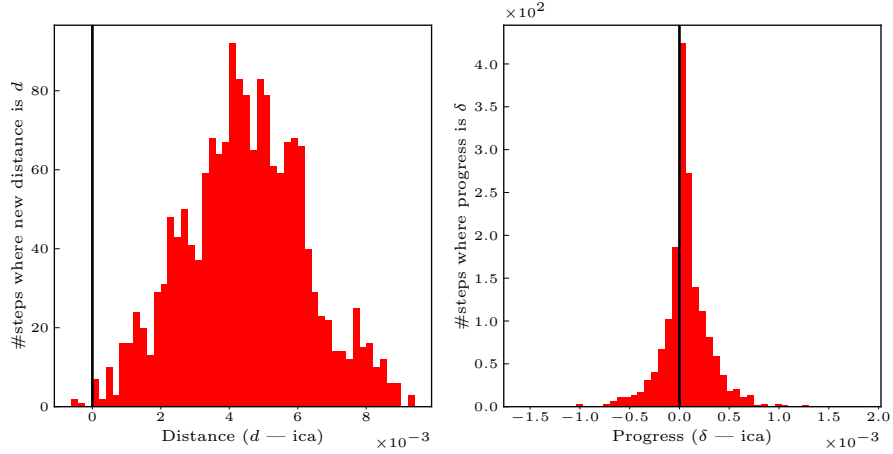

(b) with replication of values for non-targeted hidden feature components

Figure 9: Assessing the impact of constraining non-targeted hidden feature components on progress, using $\mathcal{N}_{\mathsf{ms}}$. On the left: histograms representing the distribution of all distances to the closest bounds of target hidden feature intervals during 40 runs of up to 100 iterations (with extended discretisation strategy and bins uniformly drawn from a fixed set of values) of test case generation by DeepConcolic targeting BN-based feature coverage when ICA was used as feature extraction technique—PCA gives similar-looking results, although it does not map to feature components that are comparable with each other. On the right: similar histograms showing the difference between the distance of a source test input x and a newly generated one x': $\delta > 0$ indicates that the generated input is closer to the target hidden feature interval than its source test. Contrary to the top row (a), the bottom row (b) shows distributions for d and δ obtained when the counter measure of Eq. (19) was enabled.

(a) without replication of values for non-targeted hidden feature components

(b) with replication of values for no-targeted hidden feature components

Figure 10: Assessing the impact of max-pooling layers, and constraining non-targeted hidden feature components on progress, using $\mathcal{N}_{\mathsf{mx}}$. See Figure 9 for a description of the histograms.

we briefly review a few directions that are closely related to this paper. Please refer to [16] for a comprehensive review.

Testing   includes methods to generate a set of test cases and use the generated test cases to evaluate the reliability (or other properties) of deep learning. There are a number of ways to determine how the test cases are generated, including *e.g.,* fuzzing [11], coverage metrics [34, 14], symbolic execution [10, 31], concolic testing [32], etc.

Verification   includes algorithms to check if a deep learning model satisfies a property. Existing verification algorithms can be roughly categorised into constraint-solving based methods [19], abstract interpretation based methods [9, 21, 29], global optimisation based methods [15, 27, 28], and game-based methods [38, 39].

Abstraction   There are a few existing methods on abstraction of deep learning. For example, in [6], a Boolean abstraction on the ReLU activation pattern of some specific layer is considered and monitored. Conversely of Boolean abstraction, Henzinger et al. [12] and Cheng [5] consider box abstractions. The qualitative box abstraction is enhanced by Lukina et al. [23] to identify unexpected inputs with a quantitative abstraction of box distance, which measures the minimum distance to one of the class centres.

Probabilistic Assessment   unlike testing and verification which are mainly to determine the existence of "bugs" (*i.e.,* counterexamples to the satisfiability of desirable properties) in the deep learning model, assesses the satisfiability of a property in a probabilistic way, by *e.g.,* summarising over sampling results. Due to the unknown ground truth over the underlying distribution, this approach usually requires to either make an assumption over the underlying distribution [37] or learn the distribution [41]. Note that the obtained probabilistic assessment can be bounded [3, 18].

Certification and Safety Arguments   While the above techniques may compute the evidence to the (un)satisfiability of a property, the properties they work with are mainly low-level ones, such as the point-wise robustness which only regards robustness *w.r.t.* a given input. More investigations are needed to understand if and how the evidence we can gain against these low-level specifications, can contribute to the claim of higher-level specifications such as "a deep learning model can be free from failure for the next k inferences" [40]. One step further, it might be interesting to understand the safety issues in learning-enabled autonomous systems [30, 13].

9 Conclusions

In this paper, we have suggested a new way to abstract the successive layers of a trained DNN by means of a Bayesian Network. The abstraction relation between these two objects is parameterised by a linear feature extraction mapping and a discretisation strategy for the features.

The constructed model can also be used to assess the quality of a set of test inputs *w.r.t.* both the amount of individual features it covers, and assumed relationships between features. We have exercised this ability by extending the test case generation tool DeepConcolic with our new BN-based coverage criteria. We have carried out extensive experiments to assess the effectiveness of this tool in generating new inputs despite the approximations that stem from the feature extraction techniques, as well as the linear encoding of non-linear functional behaviours (like max-pooling layers). Our experimental results show that the new tool is able to efficiently generate new test inputs, including adversarial examples, and generates new inputs by emphasising portions of learnt features to legitimate inputs. However, the formulation of our coverage metrics suffers from the discrete structure of the underlying BN, and fails to precisely account for any small-scale improvement towards fulfilling the associated criteria.

References

[1] Rob Ashmore and Matthew Hill. Boxing clever: Practical techniques for gaining insights into training data and monitoring distribution shift. In *First International Workshop on Artificial Intelligence Safety Engineering*, 2018.

[2] Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, ICML'12, page 1467–1474, Madison, WI, USA, 2012. Omnipress. ISBN 9781450312851.

[3] Peter Bishop, Robin Bloomfield, Bev Littlewood, Peter Popov, Andrey Povyakalo, and Lorenzo Strigini. A conservative bound for the probability of failure of a 1-out-of-2 protection system with one hardware-only and one software-based protection train. *Reliability Engineering & System Safety*, 130:61–68, 2014.

[4] Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 267–284, Santa Clara, CA, August 2019. USENIX Association. ISBN 978-1-939133-06-9. URL https://www.usenix.org/conference/usenixsecurity19/presentation/carlini.

[5] Chih-Hong Cheng. Provably-robust runtime monitoring of neuron activation patterns, 2020.

[6] Chih-Hong Cheng, Georg Nührenberg, and Hirotoshi Yasuoka. Runtime monitoring neuron activation patterns. *CoRR*, abs/1809.06573, 2018. URL http://arxiv.org/abs/1809.06573.

[7] Vitaly Feldman. Does learning require memorization? A short tale about a long tail. *CoRR*, abs/1906.05271, 2019. URL http://arxiv.org/abs/1906.05271.

[8] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *CCS2015*, 2015.

[9] Timon Gehr, Matthew Mirman, Dana Drachsler-Cohen, Petar Tsankov, Swarat Chaudhuri, and Martin Vechev. AI2: Safety and robustness certification of neural networks with abstract interpretation. In *Security and Privacy (SP), 2018 IEEE Symposium on*, 2018.

[10] Divya Gopinath, Kaiyuan Wang, Mengshi Zhang, Corina S. Pasareanu, and Sarfraz Khurshid. Symbolic execution for deep neural networks. *CoRR*, abs/1807.10439, 2018. URL http://arxiv.org/abs/1807.10439.

[11] Jianmin Guo, Yu Jiang, Yue Zhao, Quan Chen, and Jiaguang Sun. Dlfuzz: Differential fuzzing testing of deep learning systems. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ESEC/FSE 2018, page 739–743, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355735. doi: 10.1145/3236024. 3264835. URL https://doi.org/10.1145/3236024.3264835.

[12] Thomas A. Henzinger, Anna Lukina, and Christian Schilling. Outside the box: Abstraction-based monitoring of neural networks. *CoRR*, abs/1911.09032, 2019. URL http://arxiv.org/abs/1911.09032.

[13] W. Huang, Y. Zhou, Y. Sun, J. Sharp, S. Maskell, and X. Huang. Practical verification of neural network enabled state estimation system for robotics. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7336–7343, 2020. doi: 10.1109/IROS45743.2020. 9340720.

[14] Wei Huang, Youcheng Sun, James Sharp, Wenjie Ruan, Jie Meng, and Xiaowei Huang. Coverage guided testing for recurrent neural networks, 2020.

[15] Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. Safety verification of deep neural networks. In *CAV2017*, pages 3–29, 2017. doi: 10.1007/978-3-319-63387-9_1. URL https://doi.org/10.1007/978-3-319-63387-9_1.

[16] Xiaowei Huang, Daniel Kroening, Wenjie Ruan, James Sharp, Youcheng Sun, Emese Thamo, Min Wu, and Xinping Yi. A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability. *Computer Science Review*, 37: 100270, 2020. ISSN 1574-0137. doi: https://doi.org/10.1016/j.cosrev.2020. 100270. URL http://www.sciencedirect.com/science/article/pii/S1574013719302527.

[17] A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4):411 – 430, 2000. ISSN 0893-6080. doi: https://doi.org/10.1016/S0893-6080(00)00026-5. URL http://www.sciencedirect.com/science/article/pii/S0893608000000265.

[18] Gaojie Jin, Xinping Yi, Liang Zhang, Lijun Zhang, Sven Schewe, and Xiaowei Huang. How does weight correlation affect the generalisation ability of deep neural networks. In *NeurIPS'20*, 2020.

[19] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *CAV2017*, pages 97–117, 2017.

[20] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models*. MIT Press, 2009.

[21] Jianlin Li, Jiangchao Liu, Pengfei Yang, Liqian Chen, Xiaowei Huang, and Lijun Zhang. Analyzing deep neural networks with symbolic propagation: Towards higher precision and faster verification. In Bor-Yuh Evan Chang, editor, *Static Analysis*, pages 296–319, Cham, 2019. Springer International Publishing. ISBN 978-3-030-32304-2.

[22] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. In *Network and Distributed Systems Security (NDSS) Symposium 2018*, 2018.

[23] Anna Lukina, Christian Schilling, and Thomas A. Henzinger. Into the unknown: Active monitoring of neural networks, 2020.

[24] Lei Ma, Felix Juefei-Xu, Jiyuan Sun, Chunyang Chen, Ting Su, Fuyuan Zhang, Minhui Xue, Bo Li, Li Li, Yang Liu, Jianjun Zhao, and Yadong Wang. DeepGauge: Comprehensive and multi-granularity testing criteria for gauging the robustness of deep learning systems. In *ASE2018*, 2018.

[25] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

[26] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. DeepXplore: Automated Whitebox Testing of Deep Learning Systems. In *Proceedings of the 26th Symposium on Operating Systems Principles*, SOSP '17, page 1–18, New York, NY, USA, 2017. Association for Computing Machinery.

ISBN 9781450350853. doi: 10.1145/3132747.3132785. URL https://doi.org/10.1145/3132747.3132785.

[27] Wenjie Ruan, Xiaowei Huang, and Marta Kwiatkowska. Reachability analysis of deep neural networks with provable guarantees. In *IJCAI*, pages 2651–2659, 2018.

[28] Wenjie Ruan, Min Wu, Youcheng Sun, Xiaowei Huang, Daniel Kroening, and Marta Kwiatkowska. Global robustness evaluation of deep neural networks with provable guarantees for the hamming distance. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5944–5952. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/824. URL https://doi.org/10.24963/ijcai.2019/824.

[29] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.*, 3(POPL), January 2019. doi: 10.1145/3290354. URL https://doi.org/10.1145/3290354.

[30] Y. Sun, Y. Zhou, S. Maskell, J. Sharp, and X. Huang. Reliability validation of learning enabled vehicle tracking. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9390–9396, 2020. doi: 10.1109/ICRA40945.2020.9196932.

[31] Youcheng Sun, Xiaowei Huang, and Daniel Kroening. Testing deep neural networks. *CoRR*, abs/1803.04792, 2018. URL http://arxiv.org/abs/1803.04792.

[32] Youcheng Sun, Min Wu, Wenjie Ruan, Xiaowei Huang, Marta Kwiatkowska, and Daniel Kroening. Concolic testing for deep neural networks. In *ASE*, 2018.

[33] Youcheng Sun, Xiaowei Huang, and Daniel Kroening. Structural coverage metrics for deep neural networks. *EMSOFT2019*, 2019.

[34] Youcheng Sun, Xiaowei Huang, Daniel Kroening, James Sharp, Matthew Hill, and Rob Ashmore. Structural test coverage criteria for deep neural networks. In *Proceedings of the 41st International Conference on Software Engineering: Companion Proceedings*, ICSE '19, pages 320–321, Piscataway, NJ, USA, 2019. IEEE Press. doi: 10.1109/ICSE-Companion.2019.00134. URL https://doi.org/10.1109/ICSE-Companion.2019.00134.

[35] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR2014*, 2014.

[36] Sergios Theodoridis and Konstantinos Koutroumbas. *Pattern recognition*. Elsevier, 2003.

[37] Tsui-Wei Weng, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, Dong Su, Yupeng Gao, Cho-Jui Hsieh, and Luca Daniel. Evaluating the robustness of neural networks: An extreme value theory approach, 2018.

[38] Matthew Wicker, Xiaowei Huang, and Marta Kwiatkowska. Feature-guided black-box safety testing of deep neural networks. In Dirk Beyer and Marieke Huisman, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 408–426, Cham, 2018. Springer International Publishing. ISBN 978-3-319-89960-2.

[39] Min Wu, Matthew Wicker, Wenjie Ruan, Xiaowei Huang, and Marta Kwiatkowska. A game-based approximate verification of deep neural networks with provable guarantees. *arXiv preprint arXiv:1807.03571*, 2018.

[40] Xingyu Zhao, Alec Banks, James Sharp, Valentin Robu, David Flynn, Michael Fisher, and Xiaowei Huang. A safety framework for critical systems utilising deep neural networks. In *SafeComp'20*, volume 12234 of *LNCS*, pages 244–259, 2020.

[41] Xingyu Zhao, Wei Huang, Alec Banks, Victoria Cox, David Flynn, Sven Schewe, and Xiaowei Huang. Assessing reliability of deep learning through robustness evaluation and operational testing. In *SafeComp'21*, volume 12234 of *LNCS*, 2021.