

Outside the Box: Abstraction-Based Monitoring of Neural Networks

Thomas A. Henzinger and Anna Lukina and Christian Schilling¹

Abstract. Neural networks have demonstrated unmatched performance in a range of classification tasks. Despite numerous efforts of the research community, *novelty detection* remains one of the significant limitations of neural networks. The ability to identify previously unseen inputs as novel is crucial for our understanding of the decisions made by neural networks. At runtime, inputs not falling into any of the categories learned during training cannot be classified correctly by the neural network. Existing approaches treat the neural network as a black box and try to detect novel inputs based on the confidence of the output predictions. However, neural networks are not trained to reduce their confidence for novel inputs, which limits the effectiveness of these approaches. We propose a framework to monitor a neural network by observing the hidden layers. We employ a common abstraction from program analysis—boxes—to identify novel behaviors in the monitored layers, i.e., inputs that cause behaviors *outside the box*. For each neuron, the boxes range over the values seen in training. The framework is efficient and flexible to achieve a desired trade-off between raising false warnings and detecting novel inputs. We illustrate the performance and the robustness to variability in the unknown classes on popular image-classification benchmarks.

1 INTRODUCTION

Neural networks have become the state of the art for a wide range of academic and industrial machine-learning applications, such as image or speech recognition [44, 42, 47]. With this technology becoming ever more widespread, one of the next great challenges is building techniques for identifying and mitigating intrinsic limitations of neural networks in the general problem domain of classification. Given an input, a neural-network classifier must, by definition, output one of the classes it was trained for. The ability to output “do not know” for *novel inputs* (i.e., inputs corresponding to classes the network was not trained for) is crucial for safety-critical applications. The software architects of autonomous cars, for instance, are facing a trade-off between efficiency and risk to misclassify anomalies [6, 46, 33]. This fundamental problem of novelty detection has been of great interest to the research community (see the survey [36]). Moreover, evaluating learning algorithms in the face of parameter or input variation has become a part of the emerging topic of explainable artificial intelligence [13, 31], where interpretability is investigated as a way to ascertain reliability of a learned system.

In search of a deeper understanding of the neural network’s decision making and improved runtime management of the novel inputs, we turn to abstraction techniques commonly used in program analysis for monitoring complex safety-critical systems [11, 29, 8, 2].

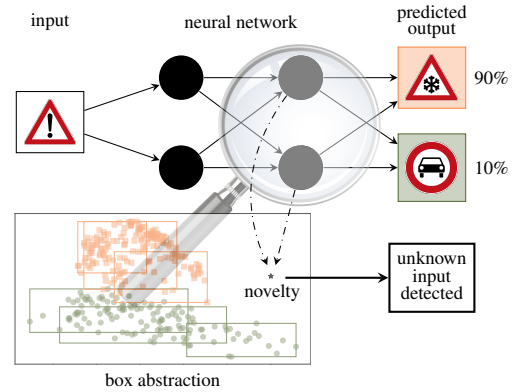


Figure 1: A schematic view of our monitoring framework.

Focusing on novelty detection, we propose to accompany the neural-network classifier with a *runtime monitor* that supervises the decisions. Fig. 1 depicts the high-level architecture of our framework. Running in parallel, the neural network and the monitor share the same interface, which allows for seamless integration into existing tools. The framework receives an input to be classified and can have two types of outputs: a classification (the neural network’s decision) or a warning (“do not know”).

While this architecture is general enough to be applied to other classification techniques, our work is built around neural networks. The monitor “watches” a number of fixed network layers chosen to incorporate the essential feature information, namely layers close to the final network output [49]. The underlying assumption is that the neurons at the watched layers exhibit a pattern typical for inputs of the same class. The monitor is trained to recognize these patterns. At runtime, the output of the watched layers is compared against the corresponding pattern. In case of close resemblance, the monitor accepts the input and the framework outputs the class predicted by the neural network. In the opposite scenario, the monitor suspects the network of making a classification decision in an atypical way. With this suspicion, the monitor rejects the proposed class and outputs a warning about a possible novelty instead.

The patterns (abstractions) we consider are (unions of) intervals, or boxes, overapproximating the set of known neuron valuations. Despite their simplicity, our experiments show a remarkable novelty-detection performance. Owing to their efficiency, boxes can be used for runtime monitoring with no significant overhead.

Our contributions can be summarized as follows. We propose an abstraction-based approach to detect novel inputs to neural-network classifiers, independent of their architecture. The abstraction at chosen layers concisely represents all values ever seen during training.

¹ IST Austria, Austria, email: {tah,anna.lukina,christian.schilling}@ist.ac.at

We can efficiently identify novel inputs at runtime by comparing the behavior of the neural network to the abstraction. Our approach can be tailored to a desired trade-off between the number of false warnings and undetected novelties.

1.1 Related work

Runtime monitoring. Runtime monitoring for machine-learned systems is a common approach in the literature. Bishop proposes to use a runtime monitor to estimate the uncertainty in a neural network, where statistical likelihood is used as the measure of novelty [5]. However, unlike boxes, computing the likelihood is expensive. The work by Gilpin considers a hierarchy of monitors: each component of a system has its own monitor, and for each subsystem consisting of several components there is a committee of monitors [17]. Dokhanchi et al. present *quality temporal logic* for specifying properties of runtime monitors about label stability in video streams [12]. Similar properties can be modeled with the *model assertions* from [22]. The abstraction in our approach could be used to explain when a label change is to be expected (namely, when, over time, the vectors observed at layer ℓ approach the border of the abstraction). Cheng et al. introduce Boolean abstraction for neural-network monitoring, which, unlike the abstraction presented in this paper, is specific to ReLU activation functions (defined as $\sigma(x) = \max(x, 0)$ for neuron x) [7]. Due to the use of operations on Boolean logic with a binary decision diagram (BDD), their approach is only scalable for layers with a few neurons.

Novelty detection. Novelty detection has been investigated by many researchers (see, e.g., [36] for a survey). It is well known that the problem stems from differences in data distributions at training and prediction time [3, 19]. Some approaches, e.g., the work by Ganin and Lempitsky [15], attempt to circumvent such cases by domain adaptation [35], which requires sampling the distribution at runtime. Other approaches try to detect novelties probabilistically, e.g., using nonparametric density estimation [23]. Few approaches, e.g., [37], perform an online adaptation of classifiers without having access to the whole distribution. Our framework is orthogonal to these stochastic techniques, since we construct an abstraction. Approaches such as *k-centers* [48] and *support-vector data description* [32] consider ideas related to the ball abstraction presented in our experimental comparison (Section 6.3), but they do not operate on neural networks. Another solution to detect novelties is *one-class classification*, where a classifier is trained to separate inputs of a single class from all other inputs (see, e.g., [38] for a recent approach).

Anomaly detection. In *selective classification* (or *abstention*) the idea is to reject inputs if a confidence score is too low [16, 1, 20, 43]. Unlike novelty detection, selective classification is applied at training time already. A well-known confidence score is the “softmax score” (see, e.g., [19, 18]), which we compare to in our experimental evaluation. Liang et al. observe that novelties may result in lower confidence scores when applying *supportive perturbation* to the input and *temperature scaling* to the output [27]. In contrast, our monitor does not require preprocessing of the input. Gal and Ghahramani illustrate limitations of the softmax output as confidence metric [14]. Their approach requires access to the network structure to add dropout functions for modeling predictive uncertainty. Lakshminarayanan et al. quantify that uncertainty using ensembles of neural networks [25].

Sun and Lampert consider a generalization of novelty called *out-of-specs situation*, which also includes the case that *situations* from




Figure 2: Left: example of a neural network that in each layer computes the function $f(\vec{x}) = \sigma(A\vec{x} + \vec{b})$ where A given by the weights, $\vec{b} = \vec{0}$, and σ is the componentwise ReLU activation function. Right: box abstraction for layer ℓ_2 and the inputs given in Ex. 1.

training never occur at runtime [41]. Although our approach targets the task of novelty detection only, according to the criteria proposed in that work, our framework is *universal* (applicable to different network architectures), *pre-trained ready* (requires no access to network training), and *nonparametric* (uses no a priori knowledge about the data distribution), but not *black-box ready* (since we require access to the network output in chosen layers).

Complementary to novelty detection, *failure prediction* is the task of finding incorrect classifications that do not arise from input novelty (see e.g., [50]). Some approaches, e.g., *open-set learning* [4] and *zero-shot learning* [34] learn new classes at training respectively at prediction time, which also requires that novel inputs can be detected.

2 PRELIMINARIES

We shortly introduce the basic terminology used in this paper.

2.1 Neural networks

We assume the reader is familiar with the basic concepts of a neural network (see, e.g., [39]). In Fig. 2 we depict a simple neural network, which we will later use to explain our approach. We do not make any assumptions about the architecture or parameters (e.g., the activation functions) of the neural network. Let \mathcal{Y} be the set of output classes. Given an n -dimensional input $\vec{x} \in \mathbb{R}^n$ and an index $\ell \in \mathbb{N}$ of a d -dimensional layer, we define the functions **watch** : $\mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}^d$ and **classify** : $\mathbb{R}^n \rightarrow \mathcal{Y}$ where **watch**(\vec{x}, ℓ) is the output at layer ℓ and **classify**(\vec{x}) is the class predicted by the network. Our approach is based on the following assumption, supported by other works [49].

Assumption 1 *The layers close to the output layer of a neural network contain high-level information. Moreover, the output at these layers is similar for inputs of the same class and different for inputs of different classes.*

2.2 Box abstraction

Our monitor will observe the output at a d -dimensional layer (i.e., it obtains vectors in \mathbb{R}^d). Given a finite set $X \subseteq \mathbb{R}^d$ of such vectors, we want to construct a set $Y \supseteq X$ that generalizes X to infinitely many elements. This concept is known as *abstraction*. The rationale is to choose a simple representation for Y that is easy to manipulate and answer queries for. In this context, we are interested in the following operations on such sets Y :

- creation of a set Y from a finite set X of vectors,
- a membership test for a vector \vec{x} (i.e., deciding $\vec{x} \in Y$), and
- (optional) enlargement (or bloating) to a superset $Y' \supseteq Y$.

Algorithm 1: Constructing abstraction at layer ℓ

Input: \mathcal{Y} : output classes
 $D = \{\langle \vec{x}^{(1)}, y^{(1)} \rangle, \dots, \langle \vec{x}^{(m)}, y^{(m)} \rangle\}$: training data
 τ : clustering parameter
Output: $A_1, \dots, A_{|\mathcal{Y}|}$: lists of abstractions

```

1 for  $y \in \mathcal{Y}$  do
    // collect all outputs at layer  $\ell$  for inputs of class  $y$ 
2    $W_y \leftarrow \left\{ \text{watch}(\vec{x}, \ell) \mid \langle \vec{x}, y \rangle \in D \wedge y = \text{classify}(\vec{x}) \right\}$ ;
3    $C_y \leftarrow \text{cluster}(W_y, \tau)$ ; // divide collected vectors into clusters
4    $A_y \leftarrow []$ ; // list of abstractions for class  $y$ 
5   for  $C \in C_y$  do
        // construct abstraction for vectors in cluster  $C$ 
6      $A_y^C \leftarrow \text{abstract}(C)$ ;
7      $A_y.\text{add}(A_y^C)$ ; // add abstraction to list
8   end
9 end
10 return  $A_1, \dots, A_{|\mathcal{Y}|}$ 

```

We focus on the *interval abstraction* [10] where the set Y is a Cartesian product of intervals $[l_i, u_i]$ with l_i and u_i being the respective lower and upper bounds in dimension i . Geometrically, the shape of Y is called a *box* (or hyperrectangle). A d -dimensional box can be represented by $2d$ bounds. Creating a tight box around a set of m vectors is a simple $\mathcal{O}(dm)$ task. Testing membership of a vector in a box is in $\mathcal{O}(d)$. Boxes can be enlarged (absolutely or relatively) by extending the bounds in $\mathcal{O}(d)$. In this work, we propose an extension to a union of such boxes, which we call the *box abstraction*.

3 OUTSIDE-THE-BOX MONITORING

In this section we describe the process of building and employing a monitor for a neural-network classifier.

3.1 Constructing a monitor

The first step is to construct a monitor for a given trained neural network and a labeled training dataset. We would typically use the same dataset that the network was trained on, but this is not obligatory. We require access to the network's output at predefined layers (which we call the *watched layers*). To simplify the presentation, we describe the concept for a single watched layer ℓ , but we discuss the generalization to multiple layers in Section 4.2.

In Algorithm 1 we present the pseudocode for constructing a monitor at layer ℓ . The algorithm consists of three phases for each output class. In the first phase (line 2), we run the network on the training data while watching layer ℓ , i.e., we collect the corresponding output at layer ℓ . The output is labeled with the corresponding ground-truth class (we only consider correctly classified input data²).

Example 1 Recall the network from Fig. 2 and the following labeled training data for output classes \blacksquare and \bullet :

$$D = \{ \langle (0.5, 0.5)^T, \blacksquare \rangle, \langle (0.5, 0.6)^T, \blacksquare \rangle, \langle (0.4, 0.6)^T, \blacksquare \rangle, \\ \langle (0.2, 0.7)^T, \blacksquare \rangle, \langle (0.7, 0.2)^T, \bullet \rangle, \langle (0.6, 0.2)^T, \bullet \rangle, \\ \langle (0.7, 0.1)^T, \bullet \rangle, \langle (0.8, 0.1)^T, \bullet \rangle, \langle (0.9, 0.2)^T, \bullet \rangle \}$$

² For the purpose of novelty detection, whether or not to consider misclassified inputs is not crucial. However, experiments suggested that ignoring misclassified inputs improves the false-negative rate of the monitor.

Algorithm 2: Monitoring at layer ℓ

Input: \vec{x} : network input
 $A_1, \dots, A_{|\mathcal{Y}|}$: lists of abstractions
Output: “accept”/“reject”: answer

```

1  $y \leftarrow \text{classify}(\vec{x})$ ; // predict class of  $\vec{x}$ 
2  $\vec{v} \leftarrow \text{watch}(\vec{x}, \ell)$ ; // collect output at layer  $\ell$ 
3 for  $A_y^C \in A_y$  do // check each abstraction for class  $y$ 
4   if  $\vec{v} \in A_y^C$  then
5     return “accept”; // found an abstraction containing  $\vec{v}$ 
6   end
7 end
8 return “reject”; //  $\vec{v}$  is not contained in any abstraction

```

Watching the second (hidden) layer ℓ_2 , we obtain the following vectors, which we label with the ground truth (also depicted in Fig. 2):

$$W_{\blacksquare} = \{(0.3, 0.45)^T, (0.38, 0.51)^T, (0.4, 0.48)^T, (0.52, 0.48)^T\}$$

$$W_{\bullet} = \{(0.02, 0.33)^T, (0.04, 0.3)^T, (0, 0.27)^T, (0, 0.3)^T, (0, 0.39)^T\}$$

Having obtained the (labeled) vectors from the training data, we continue with the second phase. Since the vectors collected for each class often cover different regions of the state space, we use a clustering algorithm to group the vectors based on their region (line 3). We note that this step is not mandatory, but we found that it can improve the precision substantially. In our implementation, we use *k-means* clustering [28], which requires to fix the number of clusters in advance; hence we iteratively increase the number of clusters until the relative improvement of the sum-of-squares metric falls below a threshold τ , which is an input parameter of the algorithm.

In the third phase, we construct a box abstraction for each combination of class and cluster identified before (function **abstract** in line 6). As a result, we obtain a list of abstractions for each class.

Example 2 Consider again Example 1. For simplicity, we assume that we obtain a single cluster for each class. For each class-cluster combination, we construct the tightest box that contains all vectors, as depicted in Fig. 2 on the right. For instance, for class \bullet the extremal values are $[0, 0.04]$ in dimension 0 and $[0.27, 0.39]$ in dimension 1, which correspond to the green box in the lower left.

3.2 Monitoring procedure

We have computed an abstraction by watching a given layer ℓ during monitor training. We now describe how the monitor uses this abstraction to operate at runtime.

Recall the architecture from Fig. 1. We summarize the pseudocode of the monitoring procedure in Algorithm 2. Given an input vector \vec{x} , we first ask the network for a prediction y (function **classify** in line 1). While the inputs are propagated through the network, we observe the output vector \vec{v} at layer ℓ (line 2). (Note that, in practice, the second step can be implemented as part of the first step instead of querying the neural network twice.) We then ask the abstractions A_y constructed for class y whether one of them contains \vec{v} (line 4). If this is the case, we conclude that the network processed the input in a usual way, and the monitor accepts the prediction y (line 5). Otherwise, we conclude that the network has processed the input in an unusual way, and the monitor rejects the input (line 8).

Example 3 In Fig. 3, we give an example of the box abstraction in an arbitrary 2D projection for one of our benchmarks. Note that the




Figure 3: Box abstraction and the outputs obtained for new data at the second-to-last layer on the MNIST benchmark (see Section 5) with two known classes (■, ●) and one unknown class (★) (projection to two arbitrary dimensions). For better visibility, we fixed the number of boxes to three per class.

boxes are created from the training dataset and the figure shows the points from the test dataset. With few exceptions, the boxes still contain the points of the same class. As can be seen in the figure, most of the novelties are not contained in any box for this projection, which leads to rejection by the monitor regardless of the class the network predicts. Assuming that the three novelties inside the projected box of class ■ are also contained in the high-dimensional box, these novelties will still be rejected if the network predicts the class ●.

4 DISCUSSION

We now discuss the detection effectiveness, the extension to multiple layers, and the resource efficiency of our monitoring approach.

4.1 Detection effectiveness

We consider the potential outcomes of a query to a monitored network. Firstly, since the abstractions *overapproximate* all vectors ever seen during training, the monitor never rejects a training input.

Proposition 1 *Given a monitored network and an input \vec{x} used for monitor training, the monitor always accepts the prediction of the network for input \vec{x} at runtime.*

Additional training always raises the monitor’s acceptance rate since new inputs can only increase the abstraction. We say that an abstraction has *converged* if it has reached a fixpoint, i.e., if additional training on any new input of a known class does not enlarge the abstraction anymore. Monitors with converged abstraction generalize the above property to any inputs of the known classes, and conversely, warnings of monitors with converged abstraction are always genuine. Consequently, false warnings can only occur for non-converged abstraction, and they indicate that more training data is required. We summarize this observation below.

Proposition 2 *Consider a monitored network with converged abstraction and an input \vec{x} . If the network predicts the correct class, the monitor accepts the prediction. If the monitor rejects the prediction, then \vec{x} does not belong to the predicted class.*

By design, an abstraction loses the direct link to the inputs it was created for. As a result, if a point is outside the abstraction, then we

can determine with certainty that this point does not belong to the set outlined by the abstraction. On the other hand, when a point lies inside the boundaries of the abstraction, we cannot conclude if the set of values inside the abstraction and the given point belong to the same input class. This property makes abstractions efficient, and determining the values outside an abstraction is sufficient for novelty detection. However, if a novelty produces layer outputs falling inside an abstraction, we would still like to reject this input. Here one could use another novelty-detection approach in parallel.

Assume the abstraction has not yet converged. We differentiate two cases of false warnings (false positives). In the first case, the vector obtained from the watched layer is outside the abstraction but close to its border. Then the abstraction is too precise, meaning that its generalization power is too weak. Generally one can enlarge the abstraction to prevent such scenarios. In the second case, the vector is far outside the abstraction. According to Assumption 1 that the output at the layer ℓ is representative of the input, it follows that the network has processed the input in a unique way never observed during training. Hence this case indicates insufficient training.

Now consider the case when the monitor accepts the input, i.e., the vector observed at layer ℓ is inside the abstraction, but the prediction of the network is wrong (false negative). With Assumption 1, the monitor should not have seen any neighboring vector during training. In this case, the abstraction is too coarse. We proposed to use clustering in order to mitigate such scenarios. As an alternative, we also experimented with more precise types of abstractions, but the results were not convincing (see Section 6.3).

We note that for each class we monitor every neuron in the layer. However, typically only a subset of the neurons (identifiable using principal component analysis [21]) is relevant for a given class. Focusing on these neurons is beyond the scope of this work.

4.2 Monitoring multiple layers

There are two main options to generalize the abstraction from one layer to multiple layers. The first option is to consider a list of monitors, one for each layer, and to accept an input if and only if all monitors accept it. The second option is to instead concatenate the outputs of the layers, i.e., for two layers ℓ_i and ℓ_j to consider $\text{watch}(\vec{x}, \ell_i) \cdot \text{watch}(\vec{x}, \ell_j)$ as the input to the (single) monitor. In this work we use the first option, which we justify as follows. While the second option is more precise as it keeps track of dependencies between neurons in different layers, the hierarchical structure of neural networks makes inter-layer dependencies less important. The second option also increases the input dimension for the monitor, which leads to the curse of dimensionality, especially for clustering.

4.3 Resource efficiency

As mentioned in Section 2, the box abstraction provides linear-time operations for creation, membership, and (relative or absolute) enlargement of each box. In particular, creating a box abstraction from samples is an incremental process, which is advantageous for processing large datasets, as we do not require to store all data at once. Note that our presentation of Algorithm 1 processes the data in one batch for clustering, but any online-clustering algorithm (e.g., [9]) can be used instead. We remark that working with a box abstraction can also be parallelized, but we use a sequential implementation.

The size of memory required to store a box is independent of the amount of training data. Thus the memory requirement is linear in the number of neurons of the watched layers and the number of clusters.

5 EXPERIMENTAL EVALUATION

In this section, we assess the performance of our monitoring framework, where we employ the following setup. To emulate the scenario of novel inputs, we train a neural network on k out of n available classes and vary k in the range $[2, n)$. We choose the first k classes in the order defined by each dataset. (Further experiments did not reveal a significant correlation between the monitoring performance and the particular ordering of the k classes.) We use the same network training parameters (e.g., number of epochs) for all instances of the same benchmark. After training the network, we construct the monitor from the same training data. To simulate convergence of the abstraction, we also include a scenario where we additionally train the box abstraction on the k classes of the test dataset. We then run the monitored network on the whole test dataset of n classes.




Since our approach is complementary to classification techniques, we analyze four popular datasets for image classification: MNIST [26], fashion MNIST (F_MNIST) [45], CIFAR-10 [24], and German traffic signs (GTSRB) [40]. We use two neural-network models from [7], which we call NN_1 and NN_2 . More information on the datasets and training parameters are provided in Tab. 1.

Table 1: Benchmarks and their parameters used in training the models from [7]. Column “Acc. train / test” shows the training and testing accuracy, respectively. Accuracy fluctuates for the models trained on different numbers of classes and we provide a range.

ID	Dataset	Classes	Inputs train / test	Network	Epochs	Acc. train / test %
1	MNIST	10	60,000 / 10,000	NN_1	10	99/99
2	F_MNIST	10	60,000 / 10,000	NN_1	30	98 – 99/91 – 99
3	CIFAR-10	10	60,000 / 10,000	NN_2	200	99/71 – 95
4	GTSRB	43	39,209 / 12,630	NN_2	10	98 – 99/88 – 97

A monitor solves a binary classification problem. Given an input and a network prediction, the task is to either accept or reject this prediction, which is also called a *negative* and a *positive* test, respectively. Novelty detection corresponds to a *true positive*. Two types of errors can occur: the input is rejected when it must be accepted (*false positive*, equivalent to a false alarm), or the input is accepted when it must be rejected (*false negative*, equivalent to a miss). In Tab. 2 we outline our notation for these cases later used in the plots.

Table 2: Notation for possible monitoring outcomes in the plots. We say that a network prediction is correct if it matches the ground truth. To save space, we do not show the true negatives; as the plots show percentages, the true negatives are the remaining difference to 100.

Pattern	Label	Explanation
	false positive	prediction correct & outside abstraction
	false negative	\neg prediction correct & \neg outside abstraction
	true positive	\neg prediction correct & outside abstraction

Our implementation and trained networks are available online.³

In the first experiment, we demonstrate the performance of our monitoring framework in comparison to two related novelty-detection approaches, which we shortly recall next.

The first approach is the *softmax prediction probability* approach following [19], which we will call the “threshold” approach. The idea




Figure 4: Softmax output for the F_MNIST benchmark with three known classes (blue (x_0), black (x_1), green (x_2)) and one unknown class (red). For each class there are 1,000 inputs.

is to interpret the softmax output of the neural network (the normalization of each output value o_i to $o_i / \sum_j o_j$) as a probability. Given a threshold value α , the approach rejects an input whenever the probability assigned to the output class falls below this threshold. We note that this approach can be restated as a special case in our framework: observing only the output layer, we use a single box for each class c_i with range $[\alpha, 1]$ in dimension i (the dimension corresponding to class c_i) and unbounded range $(-\infty, \infty)$ in all other dimensions; since α is fixed, there is no monitor training involved.

In Fig. 4, we depict the output distribution of the F_MNIST benchmark for three known and one unknown class. For three classes the outputs can in principle be distributed on a triangular plane. However, as can be seen, the network prefers to assign outputs to the corners or at least to the edges of the triangle (“+” markers in the figure), which correspond to a prediction probability 0 for two or one of the classes, respectively. For novel inputs (red dots in the figure) the network tends to do the same. Hence the threshold approach requires a high value α to reject them (we use 0.9 and 0.99 in the experiments).

We observed that the threshold approach performs poorly in scenarios with few known (or equivalently: many unknown) classes. Since for n classes the prediction probability can range in the interval $[1/n, 1]$, we normalized the threshold values relative to those intervals. For example, in an experiment with $n = 2$ known classes we map the threshold 0.9 to $[1/2, 1] \cdot 0.9 = 0.95$. For small values of n the normalized version would thus reject more inputs. Still, the threshold approach does not perform well in these scenarios.

As the second approach for comparison, we consider what we call the *Boolean abstraction* from [7]. We apply the abstraction in the last hidden layer, as suggested by the authors. Given a Boolean abstraction f , membership of a vector \vec{x} reduces to satisfiability of a Boolean formula, which is expensive to decide. Indeed, we were only able to employ the abstraction in the NN_1 network’s last hidden layer (40 neurons) and ran out of memory (8 GB) while constructing the formula on the NN_2 network’s last hidden layer (84 neurons). Compared to an average training time of 8 seconds on the 60,000 training inputs and 2 seconds for monitoring the 10,000 test inputs for our monitor in the MNIST benchmark (with four watched layers), the Boolean abstraction took 48 seconds for training and 17 seconds for running, respectively (with a single watched layer).

In Fig. 5, we compare the performance of the outlined approaches. As can be seen, learning from the test data (last bar) in addition to the training data often barely affects the novelty detection (i.e., the solid blue bars are of almost equal length) but eliminates false warnings

³ <https://github.com/VeriXAI/Outside-the-Box>




Figure 5: Comparison of novelty detection. Bars from left to right: threshold approach with $\alpha = 0.9$ and $\alpha = 0.99$, Boolean abstraction (MNIST and F_MNIST only), and box abstraction (excluding and including training on test dataset; parameters given in the captions). The x axis shows the number of known classes and the y axis shows percentages (see Tab. 2).

(outlined red bars) as expected. For the first three benchmarks we used a clustering threshold of $\tau = 0.07$. The GTSRB benchmark (for which we only report results for the instances with up to 20 classes due to lack of space) uses much fewer (less than 1,000) training data per class, which is not enough for our monitor to converge sufficiently. Hence we used a coarser value $\tau = 0.3$ and also increased the boxes by 10 % after training. Notice that for 19 known classes the number of misclassifications (the height of the solid blue and decorated yellow bars combined) produced by the network is higher than the one for the same network trained on 18 known classes, indicating that the former has inferior precision (which we validated on the test dataset). As the runtime statistics in Table 3 show, monitoring takes less than a millisecond per input despite watching multiple layers.

Table 3: Runtime statistics for different benchmarks. We report the accumulated time for clustering (“Cluster”) and for creating the abstraction (“Creation”) on all training inputs, and the time for running the monitor (“Running”) on all test inputs, averaged over all runs presented in Fig. 5. The results were obtained on a laptop (2.20 GHz CPU with four cores and 8 GB RAM) using Linux.

	MNIST	F_MNIST	CIFAR-10	GTSRB
Cluster	158.8 s	198.1 s	168.1 s	6.1 s
Creation	7.6 s	14.3 s	15.2 s	1.32 s
Running	2.1 s	4.4 s	5.6 s	9.2 s

We also experimented with combinations of our monitor and the threshold approach. For instance, we used our monitor for rejection and, in case of acceptance, used the result of the threshold approach. However, we were not able to identify a superior combination, possibly because the threshold approach is a special case of our approach.

6 FRAMEWORK FLEXIBILITY

This section is dedicated to demonstrating features of the framework that can be explored by the user to improve the desired metrics.

6.1 Monitoring multiple layers

For evaluating the combination of different layers, we monitored the MNIST benchmark with box abstraction for the last three hidden layers. We use the monitor policy to accept inputs only if the corresponding outputs at all layers are inside the respective abstraction. The detection performance is given in Fig. 6. As expected, the number of warnings increases with the number of watched layers, and we mainly see an increase in correct warnings (true positives). Interestingly, monitoring the last layer is already sufficiently precise for this benchmark on instances with many known classes (right part of the figure), but on the instances with few known classes the combination with other layers is more powerful and no layer alone can achieve the combined performance (left part of the figure).




Figure 6: Box abstraction with different observed layers. We consider all combinations of the last three layers (ℓ_x , ℓ_y , and ℓ_z) for the MNIST benchmark. The order of the bars from left to right is ℓ_z ; ℓ_y ; ℓ_x ; ℓ_y , ℓ_z ; ℓ_x , ℓ_z ; ℓ_x , ℓ_y ; ℓ_x , ℓ_y , ℓ_z .




Figure 7: Left: enlarging a box by factor γ . Right: results for enlarged boxes by different factors γ (x axis) for the MNIST benchmark with five known and unknown classes each.

6.2 Enlarging boxes

Our monitor can be easily tuned to be more (or less) restrictive by shrinking (or enlarging) the boxes. This roughly corresponds to decreasing (or increasing) the factor α in the threshold approach. We can simultaneously reason about arbitrary box sizes by computing the Euclidean distance between data points not inside any boxes and the box centers. We illustrate the idea in Fig. 7 on the left. The line segment connecting the green point with the box center has distance γd , where d is the distance between the center and the intersection of the line segment with the border of the box. Hence we would need to increase the box (i.e., each interval) by at least factor γ in order to contain the point (if we use a uniform increase factor).

Having computed the factor γ for each point, we plot the results for (uniformly) increased boxes by varying values of γ in Fig. 7 on the right. From this plot, based on the preferences, we can choose a policy to achieve, e.g., a fixed false-positive rate. For instance, to avoid false positives altogether, we should increase the boxes by factor $\gamma = 0.09$ for maximal detection. Given such a policy, the optimal value γ can be found automatically.

6.3 Comparing abstractions

Following the same main ideas as for boxes, our framework allows to consider other types of abstractions. We report results for two of




Figure 8: Evaluation of further abstractions on the MNIST benchmark. Bars from left to right: balls, octagons, and boxes in the second-to-last layer, and the bar for boxes from Fig. 5(a).

them: (Euclidean) balls and shapes called “octagons” [30]. Octagons are tighter versions of boxes with additional diagonal constraints between each pair of dimensions. Octagons can be stored as difference-bound matrices with $\mathcal{O}(d^2)$ entries, where d is the number of neurons of the watched layer. Determining whether a point lies inside a ball is linear in d (as it is for boxes), but for octagons it is quadratic. Since a ball needs to have the same radius in each dimension and neural networks do not use the same domain for each neuron, balls are too coarse to be effective. As octagon abstraction is more precise than box abstraction, octagons may detect more novelties but also raise more false warnings. We compare these abstractions to the box abstraction on the MNIST benchmark in Fig. 8, where we used the second-to-last layer (first three bars). In addition, the fourth bar shows the box abstraction from Fig. 5(a) in four layers. As can be seen, boxes in four layers strictly outperform octagons in both higher detection power and fewer false warnings, while being computationally much more efficient (2 seconds compared to 121 seconds).

7 CONCLUSIONS

Guaranteeing correctness of systems that rely on neural-network classifiers remains an important open challenge. In safety-critical applications, addressing the problem of novelty detection is crucial. The framework we propose in this paper brings us one step closer to a general methodology for developing reliable machine-learned tools. Inspired by abstraction techniques that have proved effective in the program-analysis domain, we monitor neural networks at runtime. Experimental results on common benchmarks for image classification demonstrate that our framework for constructing box abstractions of neural-network layers is effective in detecting novelties and computationally cheap. As future direction, we plan to apply our approach in a real-world setting such as monitoring neural-network controllers for cyber-physical systems.

ACKNOWLEDGMENTS

We thank Christoph Lampert and Nikolaus Mayer for fruitful discussions. This research was supported in part by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE/SHiNE) and Z211-N23 (Wittgenstein Award) and the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 754411.

REFERENCES

- [1] Amr Alexandari, Avanti Shrikumar, and Anshul Kundaje, 'A flexible and adaptive framework for abstention under class imbalance', *CoRR*, **abs/1802.07024**, (2018).
- [2] Ezio Bartocci and Yliès Falcone, *Lectures on Runtime Verification: Introductory and Advanced Topics*, volume 10457, Springer, 2018.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan, 'A theory of learning from different domains', *Machine Learning*, **79**(1-2), 151–175, (2010).
- [4] Abhijit Bendale and Terrance E. Boult, 'Towards open world recognition', in *CVPR*, pp. 1893–1902. IEEE Computer Society, (2015).
- [5] Christopher M. Bishop, 'Neural network validation: an illustration from the monitoring of multi-phase flows', in *Proceedings IEE Conference on Artificial Neural Networks*, pp. 41–45, (1993).
- [6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba, 'End to end learning for self-driving cars', *CoRR*, **abs/1604.07316**, (2016).
- [7] Chih-Hong Cheng, Georg Nührenberg, and Hirotochi Yasuoka, 'Runtime monitoring neuron activation patterns', in *DATE*, pp. 300–303, (2019).
- [8] *Handbook of Model Checking*, eds., Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, Springer, 2018.
- [9] Vincent Cohen-Addad, Benjamin Guedj, Varun Kanade, and Guy Rom, 'Online k-means clustering', *CoRR*, **abs/1909.06861**, (2019).
- [10] Patrick Cousot and Radhia Cousot, 'Static determination of dynamic properties of programs', in *International Symposium on Programming*. Dunod, (1976).
- [11] Patrick Cousot and Nicolas Halbwachs, 'Automatic discovery of linear restraints among variables of a program', in *POPL*, pp. 84–96. ACM Press, (1978).
- [12] Adel Dokhanchi, Heni Ben Amor, Jyotirmoy V. Deshmukh, and Georgios Fainekos, 'Evaluating perception systems for autonomous vehicles using quality temporal logic', in *RV*, volume 11237 of *LNCS*, pp. 409–416. Springer, (2018).
- [13] Finale Doshi-Velez and Been Kim, 'Towards a rigorous science of interpretable machine learning', *CoRR*, **abs/1702.08608**, (2017).
- [14] Yarin Gal and Zoubin Ghahramani, 'Dropout as a bayesian approximation: Representing model uncertainty in deep learning', in *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1050–1059, (2016).
- [15] Yaroslav Ganin and Victor S. Lempitsky, 'Unsupervised domain adaptation by backpropagation', in *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1180–1189, (2015).
- [16] Yonatan Geifman and Ran El-Yaniv, 'Selective classification for deep neural networks', in *NeurIPS*, pp. 4878–4887, (2017).
- [17] Leilani H. Gilpin, 'Reasonableness monitors', in *AAAI*, pp. 8014–8015. AAAI Press, (2018).
- [18] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger, 'On calibration of modern neural networks', in *ICML*, volume 70 of *PMLR*, pp. 1321–1330, (2017).
- [19] Dan Hendrycks and Kevin Gimpel, 'A baseline for detecting misclassified and out-of-distribution examples in neural networks', in *ICLR*. OpenReview.net, (2017).
- [20] Heinrich Jiang, Been Kim, Melody Y. Guan, and Maya R. Gupta, 'To trust or not to trust A classifier', in *NeurIPS*, pp. 5546–5557, (2018).
- [21] Ian T. Jolliffe, *Principal Component Analysis*, Springer Series in Statistics, Springer, 1986.
- [22] Daniel Kang, Deepti Raghavan, Peter Bailis, and Matei Zaharia, 'Model assertions for debugging machine learning', in *NeurIPS ML Sys Workshop*, (2018).
- [23] Edwin M. Knorr and Raymond T. Ng, 'A unified notion of outliers: Properties and computation', in *KDD*, pp. 219–222, (1997).
- [24] Alex Krizhevsky, 'Learning multiple layers of features from tiny images', Technical report, (2009).
- [25] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell, 'Simple and scalable predictive uncertainty estimation using deep ensembles', in *NIPS*, pp. 6402–6413, (2017).
- [26] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al., 'Gradient-based learning applied to document recognition', *Proceedings of the IEEE*, **86**(11), 2278–2324, (1998).
- [27] Shiyu Liang, Yixuan Li, and R. Srikant, 'Enhancing the reliability of out-of-distribution image detection in neural networks', in *ICLR*. OpenReview.net, (2018).
- [28] Stuart P. Lloyd, 'Least squares quantization in PCM', *Trans. Information Theory*, **28**(2), 129–136, (1982).
- [29] Antoine Miné, 'A new numerical abstract domain based on difference-bound matrices', in *PADO*, volume 2053 of *LNCS*, pp. 155–172. Springer, (2001).
- [30] Antoine Miné, 'The octagon abstract domain', *Higher-Order and Symbolic Computation*, **19**(1), 31–100, (2006).
- [31] W. James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu, 'Interpretable machine learning: definitions, methods, and applications', *CoRR*, **abs/1901.04592**, (2019).
- [32] Zineb Noumir, Paul Honeine, and Cedue Richard, 'On simple one-class classification methods', in *ISIT*, pp. 2022–2026. IEEE, (2012).
- [33] NTSB. Preliminary report released for crash involving pedestrian, Uber Technologies, Inc., test vehicle, 2018.
- [34] Mark Palatucci, Dean Pomerleau, Geoffrey E. Hinton, and Tom M. Mitchell, 'Zero-shot learning with semantic output codes', in *NIPS*, pp. 1410–1418. Curran Associates, Inc., (2009).
- [35] Sinno Jialin Pan and Qiang Yang, 'A survey on transfer learning', *Trans. Knowl. Data Eng.*, **22**(10), 1345–1359, (2010).
- [36] Marco A. F. Pimentel, David A. Clifton, Lei A. Clifton, and Lionel Tarassenko, 'A review of novelty detection', *Signal Processing*, **99**, 215–249, (2014).
- [37] Amelie Royer and Christoph H. Lampert, 'Classifier adaptation at prediction time', in *CVPR*, pp. 1401–1409, (2015).
- [38] Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli, 'Adversarially learned one-class classifier for novelty detection', in *CVPR*, pp. 3379–3388. IEEE Computer Society, (2018).
- [39] Jürgen Schmidhuber, 'Deep learning', in *Encyclopedia of Machine Learning and Data Mining*, 338–348, Springer, (2017).
- [40] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel, 'The German Traffic Sign Recognition Benchmark: A multi-class classification competition', in *IJCNN*, pp. 1453–1460, (2011).
- [41] Rémy Sun and Christoph H. Lampert, 'Ks(conf): A light-weight test if a convnet operates outside of its specifications', in *GCPR*, volume 11269 of *LNCS*, pp. 244–259. Springer, (2018).
- [42] Mingxing Tan and Quoc V. Le, 'EfficientNet: Rethinking model scaling for convolutional neural networks', in *ICML*, volume 97 of *PMLR*, pp. 6105–6114, (2019).
- [43] Sunil Thulasidasan, Tanmoy Bhattacharya, Jeff A. Bilmes, Gopinath Chennupati, and Jamal Mohd-Yusof, 'Combating label noise in deep learning using abstention', in *ICML*, volume 97 of *PMLR*, pp. 6234–6243, (2019).
- [44] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou, 'Fixing the train-test resolution discrepancy', *CoRR*, **abs/1906.06423**, (2019).
- [45] Han Xiao, Kashif Rasul, and Roland Vollgraf, 'Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms', *CoRR*, **abs/1708.07747**, (2017).
- [46] Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell, 'End-to-end learning of driving models from large-scale video datasets', in *CVPR*, pp. 3530–3538. IEEE Computer Society, (2017).
- [47] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le, 'XLNet: Generalized autoregressive pretraining for language understanding', *CoRR*, **abs/1906.08237**, (2019).
- [48] Alexander Ypma and Robert P.W. Duin, 'Support objects for domain approximation', in *ICANN*, pp. 719–724. Springer, (1998).
- [49] Matthew D. Zeiler and Rob Fergus, 'Visualizing and understanding convolutional networks', in *ECCV*, volume 8689 of *LNCS*, pp. 818–833. Springer, (2014).
- [50] Peng Zhang, Jiuling Wang, Ali Farhadi, Martial Hebert, and Devi Parikh, 'Predicting failures of vision systems', in *CVPR*, pp. 3566–3573. IEEE Computer Society, (2014).