

Provably-Robust Runtime Monitoring of Neuron Activation Patterns

Chih-Hong Cheng

DENSO AUTOMOTIVE Deutschland GmbH

Email: c.cheng@eu.denso.com

Abstract—For deep neural networks (DNNs) to be used in safety-critical autonomous driving tasks, it is desirable to monitor in operation time if the input for the DNN is similar to the data used in DNN training. While recent results in monitoring DNN activation patterns provide a sound guarantee due to building an abstraction out of the training data set, reducing false positives due to slight input perturbation has been an issue towards successfully adapting the techniques. We address this challenge by integrating formal symbolic reasoning inside the monitor construction process. The algorithm performs a sound worst-case estimate of neuron values with inputs (or features) subject to perturbation, before the abstraction function is applied to build the monitor. The provable robustness is further generalized to cases where monitoring a single neuron can use more than one bit, implying that one can record activation patterns with a fine-grained decision on the neuron value interval.

I. INTRODUCTION

For autonomous driving, perception modules implemented using deep neural networks (DNNs) are trained using data collected within the *operational design domain* (ODD). The generalization capability of machine learning implicitly implies that a well-trained DNN can only perform well on inputs *similar* to the those collected in the training data set.¹ As a consequence, for autonomous vehicles operating in an *open environment* where “black swans” undoubtedly occurs, one shall design proper monitoring mechanisms to detect if the perception system encounters any input *distant* from the training data. Upon detection, perform appropriate actions such as functional degradation or conservative maneuvers.

In this paper, we consider monitors based on building an *abstraction out of neuron activation patterns* of the training data [1], [2]. As for inputs being pixels of images, it is impossible to create a compact set representation directly on the input values, the underlying idea is to select a set of neurons to be monitored, where the monitored neurons within close-to-output layers represent high-level features. For each input to the DNN, values of all monitored neurons form a feature vector. Monitor construction is to algorithmically build a compact set representation that guarantees to contain all feature vectors (via methods such as Boolean abstraction or abstraction by min and max operators). Due to abstraction, the monitor provides a *sound* guarantee: A warning over an input implies that no input in the training data set is “close” with

the transformed feature defining the closeness. The soundness of the monitor offers a clear argument for understandability, certifiably, and homologation purposes.

While previous results in activation-based monitors have demonstrated their capability in well-engineered data sets such as MNIST or GTSRB, our deployment of the monitor in a physical laboratory setting (race track) reveals unexpected challenges. In the real-world environment, the training data is only finitely sampled from the ODD, and there is potential intrinsic randomness of the entire data collection process subject to factors such as tiny changes of lighting conditions in the day. Consequently, although the training data set is sufficient to build a perception unit, monitoring construction methods that enlarge the abstraction using the validation data set are insufficient to cure aleatory uncertainty. In other words, the monitor can generate many false alarms (i.e., the vehicle is in the ODD and nothing strange occurs, but the monitor signals a warning) to an unaccepted level (e.g., 1%) that is inappropriate to be used in autonomous driving.

Towards the above issues, we design a new monitor construction algorithm with *robustness guarantees*. The underlying idea is to replace each feature vector $f(v_{tr})$ generated by an input v_{tr} propagating through the DNN with a conservative bound $[\min_{\delta \in \Delta}(f(v_{tr} + \delta)), \max_{\delta \in \Delta}(f(v_{tr} + \delta))]$ which considers all Δ -bounded perturbation over v_{tr} . The abstraction (originally applied on the set of feature vectors) is now applied on the set of conservative bounds. The bound can be computed using state-of-the-art DNN verification methods such as boxed abstraction (interval bound propagation) [3], zonotope abstraction [4], or star sets [5]. The robustness implies that one can formally guarantee that any data point being close in the input domain (subject to Δ) can never trigger a violation in the monitor whose “closeness” is defined in the feature vector space. Finally, we also extend the robust monitor concept to a generalized monitor setup where monitoring a single neuron can use more than one bit. The multi-bit setup implies that it can record activation patterns with a fine-grained decision on the neuron value interval.

The rest of the paper is structured as follows. Section II highlights the related work and provides a general comparison over other approaches. Section III presents the basic formulation of DNN, algorithms for constructing provably robust monitors, and extensions to monitors beyond single-bit on-off activation. Finally in Section IV, we summarize our findings in the experiment and conclude with future directions.

This research is part of FOCETA project that has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 956123.

¹In layman words, DNNs are good at interpolation but bad at extrapolation.

II. RELATED WORK

The problem of understanding uncertainty in the decision making of a neural network has been extensively studied. Bayesian approaches [6], [7] learn a probability distribution on the weights of a neural network. Monte-Carlo approaches use an ensemble of networks to estimate the uncertainty; the ensemble can be generated either by differently trained networks [8] or by performing drop-out in operation [9]. The ensemble nature of Monte-Carlo approaches implies that it is costly to execute. The complementing work on monitoring neuron activation patterns [1], [2], [10], [11] can derive formal guarantees as it uses sound abstraction over the training data set. Both qualitative decisions [1], [2] and quantitative decisions [11] are made possible, and it is also used in assume-guarantee based reasoning for safety verification of highway autonomous driving systems [10]. Our work complements current results in abstraction-based neuron monitoring and can be integrated into these work with ease; it adds a new dimension of integrating symbolic reasoning inside the monitor construction process, thereby providing robustness guarantees.

Finally, the technique used in this paper is highly related to the work on verification of neural networks using abstract interpretation (e.g., [3]–[5]), where the symbolic reasoning is used not in the engineering of neural networks but rather in the construction of monitors to be used in operation.

III. ROBUST NEURON ACTIVATION PATTERN MONITORS

A. DNN and Runtime Monitors

A feed-forward DNN model consists of an input layer, multiple hidden layers and an output layer. We consider only DNNs after training, i.e., all model parameters (e.g., weights, bias) are fixed. With fixed parameters, a DNN model is a function $\mathcal{G} : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^{d_n}$, where d_0 is the input dimension and d_n is the output dimension. The model \mathcal{G} is built out of a sequence of functions $\mathcal{G} := g^n \otimes g^{n-1} \otimes \dots \otimes g^2 \otimes g^1$, where n is the number of layers in the DNN, and for $k \in \{1, \dots, n\}$, $g^k : \mathbb{R}^{d_{k-1}} \rightarrow \mathbb{R}^{d_k}$ is the transformation in the k -th layer. Given an input $v_{tr} \in \mathbb{R}^{d_0}$, the computation of \mathcal{G} is $\mathcal{G}^n(v_{tr}) = g^n(g^{n-1}(\dots(g^2(g^1(v_{tr}))))\dots)$, i.e., to perform functional composition over g^1, \dots, g^n .

For the ease of explanation, here we only formulate the simplistic monitoring setup: the abstraction-based monitor applies to a specific layer, all neurons in the layer are monitored, the decision is qualitative, and there is no additional abstraction enlargement by using the validation set. Extensions such as configuring to multi-layer monitoring as well as selecting a subset of neurons to be monitored are straightforward. We use \mathcal{G}^k to abbreviate $g^k \otimes g^{k-1} \otimes \dots \otimes g^2 \otimes g^1$, and use $\mathcal{G}^{l \rightarrow k}$ to abbreviate $g^k \otimes g^{k-1} \otimes \dots \otimes g^{l+1} \otimes g^l$. For the sake of simplicity in the established theorem, we abuse the notation and may use $\mathcal{G}^0(v_{tr})$ to represent v_{tr} (i.e., no transformation applied). Lastly, for vector v and function f that outputs a vector, we use subscripts v_j and f_j to indicate the j -th value.

Let $\mathcal{D}_{tr} \subset \mathbb{R}^{d_0}$ be the training data set, then a generic algorithm for building a *neuron activation pattern monitor* $\mathcal{M}^{(\mathcal{G},k)}$ is as follows:

```

 $\mathcal{M}^{(\mathcal{G},k)} \leftarrow \mathcal{M}_0$ 
for  $v_{tr} \in \mathcal{D}_{tr}$  do
   $\mathcal{M}^{(\mathcal{G},k)} \leftarrow \mathcal{M}^{(\mathcal{G},k)} \uplus (\text{ab}(\mathcal{G}^k(v_{tr})))$ 
end for

```

Different abstraction techniques may be applied, by setting different initial value \mathcal{M}_0 , operator \uplus and function ab .

- For monitoring the minimum and maximum values [2], [10], $\mathcal{M}^{(\mathcal{G},k)}$ equals $\langle (L_1, U_1), \dots, (L_{d_k}, U_{d_k}) \rangle$ where for $j \in \{1, \dots, d_k\}$, the pair (L_j, U_j) tracks for the complete training data set, the minimum and the maximum value of the j -th neuron. Precisely,

- \mathcal{M}_0 equals $\langle (\infty, -\infty), \dots, (\infty, -\infty) \rangle$,
- ab is the identity function, and
- $\mathcal{M}^{(\mathcal{G},k)} \uplus (v_1, \dots, v_{d_k}) := \langle (\min(L_1, v_1), \max(U_1, v_1)), \dots, (\min(L_{d_k}, v_{d_k}), \max(U_{d_k}, v_{d_k})) \rangle$.

In operation, given an input $v_{op} \in \mathbb{R}^{d_0}$, the monitor $\mathcal{M}^{(\mathcal{G},k)}$ raises a warning, denoted as $\mathcal{M}^{(\mathcal{G},k)}(v_{op}) = \text{true}$, if $\exists j \in \{1, \dots, d_k\}$ such that $\mathcal{G}_j^k(v_{op}) < L_j$ or $\mathcal{G}_j^k(v_{op}) > U_j$.

- For monitoring on-off activation patterns introduced in [1], the construction performs Boolean abstraction based on a pre-defined threshold (e.g., sign of the neuron value, or average of all visited values) to build a Boolean word. The set of Boolean words can be efficiently stored using *Binary Decision Diagrams (BDDs)* [12].

- $\mathcal{M}_0 = \emptyset$.
- Let c_1, c_2, \dots, c_{d_k} be predefined constants. The function $\text{ab}(v_1, \dots, v_{d_k})$ returns (b_1, \dots, b_{d_k}) where for $j \in \{1, \dots, d_k\}$,
 - * $b_j = 1$ if $v_j > c_j$, and
 - * $b_j = 0$ otherwise.
- $\mathcal{M}^{(\mathcal{G},k)} \uplus (b_1, \dots, b_{d_k}) := \mathcal{M}^{(\mathcal{G},k)} \cup \{(b_1, \dots, b_{d_k})\}$.

Given an input $v_{op} \in \mathbb{R}^{d_0}$, the monitor raises a warning, denoted as $\mathcal{M}^{(\mathcal{G},k)}(v_{op}) = \text{true}$, if $\text{ab}(\mathcal{G}^k(v_{op})) \notin \mathcal{M}$.

B. Robust Monitors

To build robust monitors, we first define feature-level perturbation to capture the effect on the monitored layer via a worst-case estimate.

Definition 1. Given $v_{tr} \in \mathbb{R}^{d_0}$ and an integer k_p where $0 \leq k_p < k$, define the perturbation estimate $\text{pe}_k^{\mathcal{G}}(v_{tr}, k_p, \Delta) = \langle (l_1, u_1), \dots, (l_{d_k}, u_{d_k}) \rangle$ which guarantees the following property: For $\delta_1, \dots, \delta_{k_p} \in \mathbb{R}$ where $|\delta_1|, \dots, |\delta_{k_p}| \leq \Delta$, let $\check{v} := (\check{v}_1, \dots, \check{v}_{k_p}) \in \mathbb{R}^{d_{k_p}}$ where $\forall j \in \{1, \dots, d_k\} : \mathcal{G}_j^{k_p}(v_{tr}) + \delta_j = \check{v}_j$. Then the following condition holds.

$$\forall j \in \{1, \dots, d_k\} : l_j \leq \mathcal{G}_j^{k_p+1 \rightarrow k}(\check{v}) \leq u_j \quad (1)$$

Intuitively, the estimate considers perturbation occurring at the output of layer k_p , with the amount of perturbation in each dimension being bounded by Δ . When $k_p = 0$, the perturbation appears at the input layer. Any perturbed vector \check{v} , after passing through the sub-network $g^k \otimes g^{k-1} \otimes \dots \otimes g^{k_p+1}$, shall produce a vector $\mathcal{G}^{k_p+1 \rightarrow k}(\check{v})$ with every dimension being

bounded by $[l_j, u_j]$ as provided by the perturbation estimate. The perturbation estimate can be efficiently computed using boxed abstraction (interval bound propagation) [3], zonotope abstraction [4], or star sets [5]. In our implementation, boxed abstraction is used.

Let $\mathcal{D}_{tr} \subset \mathbb{R}^{d_0}$ be the training data set, then a generic algorithm for building a *robust neuron activation pattern monitor* $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}$ is as follows:

```

 $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)} \leftarrow \mathcal{M}_0$ 
for  $v_{tr} \in \mathcal{D}_{tr}$  do
   $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)} \leftarrow \mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)} \uplus_R \text{ab}_R(\text{pe}_k^{\mathcal{G}}(v_{tr}, k_p, \Delta))$ 
end for

```

We again consider previously stated two types of monitors (min-max and Boolean abstraction), by detailing the corresponding initial values \mathcal{M}_0 , operator \uplus_R and function ab_R that have different operations due to inputs being bounds.

- For robust monitoring the minimum and maximum visited neuron values, $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}$ equals $\langle (L_1, U_1), \dots, (L_{d_k}, U_{d_k}) \rangle$.
 - \mathcal{M}_0 equals $\langle (\infty, -\infty), \dots, (\infty, -\infty) \rangle$,
 - ab_R is the identity function, and
 - $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)} \uplus_R \langle (l_1, u_1), \dots, (l_{d_k}, u_{d_k}) \rangle := \langle (\min(L_1, l_1), \max(U_1, u_1)), \dots, (\min(L_{d_k}, l_{d_k}), \max(U_{d_k}, u_{d_k})) \rangle$.

In other words, the monitor considers the worst case estimate (under perturbation) and enlarges the bound correspondingly.

- For robust monitoring on-off activation patterns,
 - $\mathcal{M}_0 = \emptyset$.
 - Let c_1, c_2, \dots, c_{d_k} be predefined constants. The function $\text{ab}_R(\langle (l_1, u_1), \dots, (l_{d_k}, u_{d_k}) \rangle)$ returns (b_1, \dots, b_{d_k}) where for $j \in \{1, \dots, d_k\}$,
 - * set $b_j = 1$ if $l_j > c_j$,
 - * set $b_j = 0$ if $u_j \leq c_j$,
 - * otherwise, set b_j to “-”, where symbol “-” means “don’t care”, i.e., the value can be either 1 or 0.
 - $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)} \uplus_R (b_1, \dots, b_{d_k}) := \mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)} \cup \text{word2set}(b_1, \dots, b_{d_k})$, where $\text{word2set}()$ is the function that returns the set of all binary words by translating “-” into 0 and 1. Note that when using BDDs, the translation $\text{word2set}()$ does not create an exponential blow-up.²

Given an input $v_{op} \in \mathbb{R}^{d_0}$ encountered in operation time, as there is no perturbation (i.e., $\Delta = 0$), it implies that for every $j \in \{1, \dots, d_k\}$, the computed lower-bound l_j , the upper-bound u_j and $\mathcal{G}_j^k(v_{op})$ are equal. One can simply use the condition as stated in the standard monitor to perform abnormality detection.

- For min-max monitors, given an input $v_{op} \in \mathbb{R}^{d_0}$, the monitor $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}$ raises a warning, denoted as $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}(v_{op}) = \text{true}$, if $\exists j \in \{1, \dots, d_k\}$ such that $\mathcal{G}_j^k(v_{op}) < L_j$ or $\mathcal{G}_j^k(v_{op}) > U_j$.

²Consider the case where $(b_1, b_2, b_3, b_4) = (1, -, -, 0)$. When using BDDs, the set is simply represented as $b_1 \wedge \neg b_4$, i.e., both b_3 and b_4 are unconstrained.

- For Boolean abstraction monitors, given an input $v_{op} \in \mathbb{R}^{d_0}$, the monitor raises a warning, denoted as $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}(v_{op}) = \text{true}$, if $\text{ab}(\mathcal{G}^k(v_{op})) \notin \mathcal{M}$.

The abstraction kept by the robust monitor, upon construction, already accounts for all possible cases caused by perturbation bounded by Δ at layer k_p . Therefore, for any input similar to the training data, where the similarity is defined using Δ and k_p , the robust monitor can guarantee not to raise a warning.

Lemma 1. *Given a robust monitor $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}$, for any input $v_{op} \in \mathbb{R}^{d_0}$, if $\mathcal{M}^{(\mathcal{G}, k, k_p, \Delta)}(v_{op}) = \text{true}$, then there **does not exist** an input $v_{tr} \in \mathcal{D}_{tr}$ s.t. the following condition holds:*

$$\forall j \in \{1, \dots, d_{k_p}\} : |\mathcal{G}_j^{k_p}(v_{op}) - \mathcal{G}_j^{k_p}(v_{tr})| \leq \Delta \quad (2)$$

C. Robust Neuron Interval Activation Monitors

Finally, we propose a generic setting as an extension of simple min-max and on-off activation monitors. In contrast to only using one bit to monitor a neuron, one can use multiple bits to encode the *interval* where the neuron value is located. For ease of explanation, we use an example of 2 bits. For $j \in \{1, \dots, d_k\}$, let c_{j1}, c_{j2}, c_{j3} be predefined constants where $c_{j1} < c_{j2} < c_{j3}$.

1) *Standard Interval Activation Monitors:* For standard monitors, we have the following formulation.

- $\mathcal{M}_0 = \emptyset$.
- The function $\text{ab}(v_1, \dots, v_{d_k})$ returns (b_1, \dots, b_{d_k}) where for $j \in \{1, \dots, d_k\}$,
 - $b_j = 11$ if $v_j > c_{j3}$,
 - $b_j = 10$ if $c_{j3} \geq v_j \geq c_{j2}$,
 - $b_j = 01$ if $c_{j2} > v_j > c_{j1}$,
 - $b_j = 00$ otherwise.

In other words, for the interval activation monitor, we allow the monitor to use two bits to encode the value.³

- $\mathcal{M}^{(\mathcal{G}, k)} \uplus (b_1, \dots, b_{d_k}) := \mathcal{M}^{(\mathcal{G}, k)} \cup \{(b_1, \dots, b_{d_k})\}$.

Given an input $v_{op} \in \mathbb{R}^{d_0}$, the monitor raises a warning, denoted as $\mathcal{M}^{(\mathcal{G}, k)}(v_{op}) = \text{true}$, if $\text{ab}(\mathcal{G}^k(v_{op})) \notin \mathcal{M}$.

2) *Robust Interval Activation Monitors:* The following formulation extends the concept to a robust monitoring setting.

- $\mathcal{M}_0 = \emptyset$,
- The function $\text{ab}_R(\langle (l_1, u_1), \dots, (l_{d_k}, u_{d_k}) \rangle)$ returns (b_1, \dots, b_{d_k}) where for $j \in \{1, \dots, d_k\}$,
 - $b_j = \{11\}$ if $l_j > c_{j3}$,
 - $b_j = \{10\}$ if $c_{j3} \geq u_j \geq l_j \geq c_{j2}$,
 - $b_j = \{01\}$ if $c_{j2} > u_j \geq l_j > c_{j1}$,
 - $b_j = \{00\}$ if $c_{j1} \geq u_j$,
 - $b_j = \{00, 01\}$ if $c_{j2} > u_j > c_{j1}$ and $c_{j1} \geq l_j$,
 - $b_j = \{01, 10\}$ if $c_{j3} \geq u_j \geq c_{j2}$ and $c_{j2} > l_j > c_{j1}$,
 - $b_j = \{10, 11\}$ if $u_j > c_{j3}$ and $c_{j3} \geq l_j \geq c_{j2}$,
 - $b_j = \{00, 01, 10\}$ if $c_{j1} \geq l_j$ and $c_{j3} \geq u_j \geq c_{j2}$.

³It is a generalization of min-max monitor, as one can set c_{j3} to be the maximum visited neuron value, c_{j2} to be the minimum visited neuron value, and $c_{j1} = -\infty$. It is also a generalization of on-off monitor, by setting $c_{j3} = \infty$ and $c_{j1} = -\infty$.




Fig. 1. Understanding the encoding for robust interval activation monitors




Fig. 2. Experimenting monitoring techniques in the lab setting.

- $b_j = \{01, 10, 11\}$ if $u_j > c_{j3}$ and $c_{j2} > l_j > c_{j1}$,
- otherwise, $b_j = \{00, 01, 10, 11\}$.

In contrast to the standard monitor where the value v_j can only fall into one interval, the robust monitor maintains the conservative lower-bound l_j and upper-bound u_j , and it is possible that l_j and u_j do not fall into the same interval, as illustrated in Figure 1.

- $\mathcal{M}^{(g,k,k_p,\Delta)} \uplus_R (b_1, \dots, b_{d_k}) := \mathcal{M}^{(g,k,k_p,\Delta)} \cup \text{word2set}(b_1, \dots, b_{d_k})$, where $\text{word2set}()$ is the function that returns the set of all binary words

$$\{(\hat{b}_1, \dots, \hat{b}_{d_k}) \mid \forall i \in \{1, \dots, d_k\} : \hat{b}_i \in b_i\}$$

The function word2set can be efficiently implemented in BDD, and here we only explain the process using an example where $(b_1, b_2, b_3) = (\{00, 01\}, \{01, 10, 11\}, \{10\})$. For b_1 (similarly for b_2 and b_3), one uses two BDD variables b_{10}, b_{11} to encode the first and the second bit. Therefore, the BDD encoding of $\text{word2set}(b_1, b_2, b_3)$ simply returns $(\neg b_{10}) \wedge (b_{20} \vee b_{21}) \wedge (b_{30} \wedge \neg b_{31})$.

IV. CONCLUDING REMARKS

In this paper, we addressed the robustness of activation-pattern monitoring and proposed integrating the symbolic reasoning in the monitor construction process. As established in the new monitor construction algorithm, the provable robustness claims ensure that any input data being similar (subject to k_p and Δ) to the training data set will not cause the monitor to raise a warning. Therefore, the new monitor reduces false positives and only focuses on inputs that are genuinely distant from the training data set. We have also extended the robust monitor concept to a generic setting where each neuron can be monitored using an activation pattern expressed using more than one bit. With a particular focus on operation monitoring, this work well complements existing results that also apply formal methods to improve the quality of DNNs.

As a preliminary evaluation, we implemented the proposed technology into a lab setting, where we engineered a DNN to generate visual waypoints from images and deployed the monitor to detect abnormal situations. For the optimal case with 0.62% of false positives (i.e., the vehicle is inside ODD but warning is raised), the monitor can detect (with a substantially higher rate) our additionally created scenarios that are not included in the training data set, such as dark conditions, construction site, or ice on the track (illustrated in Figure 2). By changing to robust monitors, the false positive rate is further reduced to 0.125% (i.e., 80% reduction) while the detection rate of ODD departures remains roughly the same. However, we also discovered that some monitors, although demonstrating 0% false positive, are inefficient in that only a few warnings are raised. We expect to conduct further studies to understand how to train networks with better monitorability and to evaluate the technique in real-world scenarios such as multi-view 3D reconstruction [13], [14]. This helps to understand the full potential and limitations of the approach. We also plan to integrate the correctness claims into recently developed results in logical safety argumentation [15], [16].

REFERENCES

- [1] C.-H. Cheng, G. Nührenberg, and H. Yasuoka, “Runtime monitoring neuron activation patterns,” in *DATE*. IEEE, 2019, pp. 300–303.
- [2] T. A. Henzinger, A. Lukina, and C. Schilling, “Outside the box: Abstraction-based monitoring of neural networks,” in *ECAI*. IOS Press, 2020, pp. 2433–2440.
- [3] S. Goyal, K. Dvijotham, R. Stanforth, R. Bunel, C. Qin, J. Uesato, S. Arandjelovic, T. Mann, and P. Kohli, “On the effectiveness of interval bound propagation for training verifiably robust models,” *arXiv preprint 1810.12715*, 2018.
- [4] T. Gehr, M. Mirman, D. Drachler-Cohen, P. Tsankov, S. Chaudhuri, and M. Vechev, “Ai2: Safety and robustness certification of neural networks with abstract interpretation,” in *S&P (Oakland)*. IEEE, 2018, pp. 3–18.
- [5] H.-D. Tran, D. M. Lopez, P. Musau, X. Yang, L. V. Nguyen, W. Xiang, and T. T. Johnson, “Star-based reachability analysis of deep neural networks,” in *FM*. Springer, 2019, pp. 670–686.
- [6] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, “Weight uncertainty in neural network,” in *ICML*, 2015, pp. 1613–1622.
- [7] J. M. Hernández-Lobato and R. Adams, “Probabilistic backpropagation for scalable learning of bayesian neural networks,” in *ICML*, 2015, pp. 1861–1869.
- [8] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *NeurIPS*, 2017, pp. 6402–6413.
- [9] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *ICML*, 2016, pp. 1050–1059.
- [10] C.-H. Cheng, C.-H. Huang, T. Brunner, and V. Hashemi, “Towards safety verification of direct perception neural networks,” in *DATE*. IEEE, 2020, pp. 1640–1643.
- [11] A. Lukina, C. Schilling, and T. A. Henzinger, “Into the unknown: Active monitoring of neural networks,” *arXiv preprint 2009.06429*, 2020.
- [12] R. E. Bryant, “Symbolic boolean manipulation with ordered binary-decision diagrams,” *ACM Computing Surveys (CSUR)*, vol. 24, no. 3, pp. 293–318, 1992.
- [13] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, “Multi-view 3d object detection network for autonomous driving,” in *CVPR*, 2017, pp. 1907–1915.
- [14] B. Yang, W. Luo, and R. Urtasun, “Pixor: Real-time 3d object detection from point clouds,” in *CVPR*, 2018, pp. 7652–7660.
- [15] T. A. Beyene and A. Sahu, “Rule-based safety evidence for neural networks,” in *WAISE*. Springer, 2020, pp. 328–335.
- [16] X. Zhao, A. Banks, J. Sharp, V. Robu, D. Flynn, M. Fisher, and X. Huang, “A safety framework for critical systems utilising deep neural networks,” in *SafeComp*. Springer, 2020, pp. 244–259.