

Towards Safety Verification of Direct Perception Neural Networks

Chih-Hong Cheng^{*}, Chung-Hao Huang[†], Thomas Brunner[†] and Vahid Hashemi[‡]

^{*}DENSO AUTOMOTIVE Deutschland GmbH

[†]fortiss - Research Institute of the Free State of Bavaria

[‡]Audi AG

Contact: c.cheng@denso-auto.de

Abstract—We study the problem of safety verification of direct perception neural networks, where camera images are used as inputs to produce high-level features for autonomous vehicles to make control decisions. Formal verification of direct perception neural networks is extremely challenging, as it is difficult to formulate the specification that requires characterizing input as constraints, while the number of neurons in such a network can reach millions. We approach the specification problem by learning an input property characterizer which carefully extends a direct perception neural network at close-to-output layers, and address the scalability problem by a novel assume-guarantee based verification approach. The presented workflow is used to understand a direct perception neural network (developed by Audi) which computes the next waypoint and orientation for autonomous vehicles to follow.

Index Terms—formal verification, neural network, dependability, autonomous driving

I. INTRODUCTION

Using deep neural networks has been the *de facto* choice for developing visual object detection function in automated driving. Nevertheless, in the autonomous driving workflow, the neural networks can also be used more extensively. An example is *direct perception* [2]; one trains a neural network to read high-dimensional inputs (such as images from camera or point clouds from lidar) and produce low-dimensional information called *affordances* (e.g., safe maneuver regions or the next waypoint to follow) which could be used to program a controller for the autonomous vehicle. One may use direct perception as a hot standby system for a classical *mediated perception* system that extracts objects and identifies lane markings before affordances are produced.

In this paper, we study the *safety verification problem* for a neural network implementing direct perception, where the goal is to ensure that under certain input conditions, the undesired output values never occur. An example of such a kind can be the following: “For every input image where the road in the image strongly bends to the right, the output of the neural network should never suggest to strongly steer to the left”. Overall, the safety verification problem for direct perception networks is fundamentally challenging due to two factors:

- **(Specification)** To perform safety verification, one premise is to have the undesired property formally specified. Nevertheless, it is practically impossible to characterize input specifications from images such as

“road strongly bends to the right” and represent them as constraints over input variables.

- **(Scalability)** Neural networks for direct perception often take images with millions of pixels, and the internal structure of the network can have many layers. This challenges any state-of-the-art formal analysis framework in terms of scalability.

Towards these issues, we present a workflow for safety verification of direct perception neural networks by simultaneously addressing the specification and the scalability problem. For the ease of understanding, we use Figure 1 to explain the concept. First, we address the specification problem by learning an *input property characterizer network*, where the input of the network is connected to close-to-output layer neurons of the original direct perception network. In Figure 1, the input property characterizer takes output values from the neurons n_1^{17} , n_2^{17} , n_3^{17} , n_4^{17} and n_5^{17} in the original deep perception network. For the previously mentioned specification, the input property characterizer outputs true if an input image has “road strongly bending to the right”. By doing so, the characterization of input features is aggregated to an output of a neural network. Subsequently, the safety verification problem is approached by asking if it is possible for the input-characterizing network to output true, but the output of the direct perception network demonstrates undesired values. As both the deep perception network and the input-characterizing network have shared neuron values, safety verification can be approached by only verifying close-to-output layers without losing soundness. In Figure 1, safety verification only analyzes the sub-network colored grayed, and examines if any assignment of n_1^{17} , n_2^{17} , n_3^{17} , n_4^{17} , and n_5^{17} leads to undesired output. The bounds of the neurons n_1^{17} , n_2^{17} , n_3^{17} , n_4^{17} , and n_5^{17} can be decided by static analysis (which guarantees an overly conservative bound). However, using such a bound allows to have input images that are not possible to be seen in the operating design domain (ODD)¹. Thus we are advocating an alternative

¹In training neural networks, the value for each pixel in an image is commonly re-scaled such that the re-scaled value is in the interval $[0, 1]$. Starting verification using an input domain of $[0, 1]^{d_{i_0}}$ with d_{i_0} being the number of input image pixels, the result of formal verification always creates counter-examples in formal verification where counter-example images are so distant from what can be observed in practice (such as images without textures) and are rejected by experts.

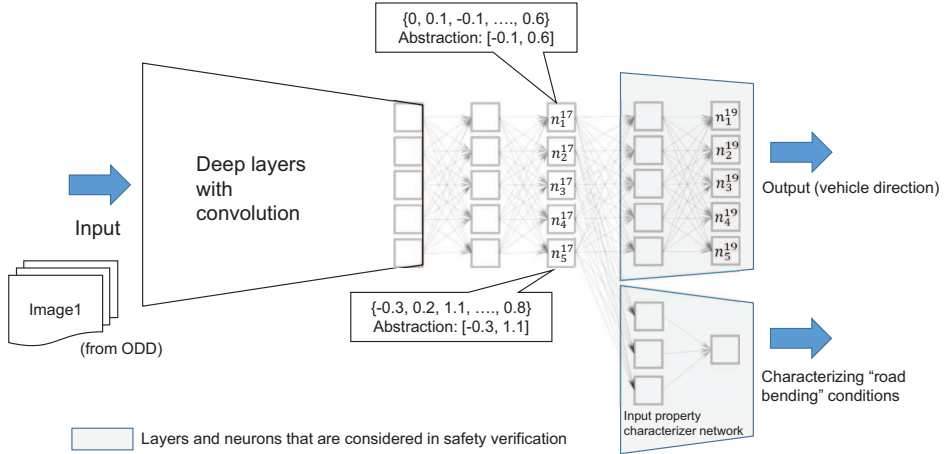


Fig. 1. High-level illustration how to perform safety verification while tackling specification and scalability issues.

assume-guarantee based approach where one first creates an outer polyhedron by aggregating all visited neuron values computed by the training set. We use the created polyhedron as a starting point to perform formal verification, by *assuming* that for every possible input data in the ODD, the computed neuron activation pattern is contained in this polyhedron. The assumption thus requires to be *monitored in runtime* by checking if any computed neuron value falls outside the polyhedron. As an example, we consider the bound of n_1^{17} to be used in verification in Figure 1. By observing the minimum and the maximum of all visited values $\{0, 0.1, -0.1, \dots, 0.6\}$, $[-0.1, 0.6]$ is an over-approximation over all visited values, and one shall monitor in runtime whether the computed value of n_1^{17} has fallen outside $[-0.1, 0.6]$.²

The rest of the paper is organized as follows. Section II presents the required definitions as well as the workflow for verification. Section III discusses extensions to a statistical setup when the input property characterizer is not perfect. Lastly, we summarize related work in Section IV and conclude with our preliminary evaluation in Section V.

II. VERIFICATION WORKFLOW

A deep neural network is comprised of L layers where operationally, the l -th layer for $l \in \{1, \dots, L\}$ of the network is a function $g^{(l)} : \mathbb{R}^{d_{l-1}} \rightarrow \mathbb{R}^{d_l}$, with d_l being the dimension of layer l . Given an input $in \in \mathbb{R}^{d_0}$, the output

²Note that such a monitoring is needed regardless of formal verification, as neurons at close-to-output layers represent high-level features, so an image in operation that leads to unexpectedly high or low neuron feature intensity (indicated by falling outside the monitored interval) can be hints for incomplete data collection or indicators for the system stepping out from the ODD.

of the l -th layer of the neural network $f^{(l)}$ is given by the functional composition of the l -th layer and the previous layers $f^{(l)}(in) := \circ_{i=1}^{(l)} g^{(i)}(in) = g^{(l)}(g^{(l-1)} \dots g^{(2)}(g^{(1)}(in)))$.

A. Characterizing Input Specification from Examples

Let $In_\phi \subseteq \mathbb{R}^{d_0}$ be the set of inputs of a neural network that satisfies the property ϕ . We assume that both ϕ and In_ϕ are unknown (e.g., the road is bending left in an image), but there exists an oracle (e.g., human) that can answer for a given input $in \in \mathbb{R}^{d_0}$, whether $in \in In_\phi$.

Let (In, C_ϕ) be the list of training data and their associated labels (generated by the oracle) related to the input property ϕ , where for every $(in, c) \in (In, C_\phi)$, $in \in \mathbb{R}^{d_0}$, $c \in \{0, 1\}$, we have $(in, 1) \in (In, C_\phi)$ if $in \in In_\phi$ and $(in, 0) \in (In, C_\phi)$ if $in \notin In_\phi$. The *perfect input property characterizer* extending the l -th layer is a function h_l^ϕ which guarantees that for every $(in, c) \in (In, C_\phi)$, $h_l^\phi(f^{(l)}(in)) = c$. The generation of h_l^ϕ can be done by training a neural network as a binary classifier, with 100% success rate on the training data. The following assumption states that as long as function h_l^ϕ performs perfectly on the training data, h_l^ϕ will also perfectly generalize to the complete input space. In other words, we can use h_l^ϕ to characterize ϕ .

Assumption 1 (Perfect Generalization). *Assume that h_l^ϕ also perfectly characterizes ϕ , i.e., $\forall in \in \mathbb{R}^{d_0}$: $h_l^\phi(f^{(l)}(in)) = 1$ iff $in \in In_\phi$.*

Definition 1 (Safety Verification). *The safety verification problem asks if there exists an input $in \in In_\phi$ such that $f^{(L)}(in)$ satisfies ψ , where the **risk condition** ψ is a conjunction of linear inequalities over the output of the neural network. If no such input in exists, we say that the neural network is safe under the input constraint ϕ and the output risk constraint ψ .*

When Assumption 1 holds, for safety verification it is equivalent to ask whether there exists an input $\text{in} \in \mathbb{R}^{d_0}$ such that $h_i^\phi(f^{(l)}(\text{in})) = 1$ and $f^{(L)}(\text{in})$ satisfies ψ . From now on, unless explicitly specified, we consider only situations where Assumption 1 holds.

B. Practical Safety Verification

a) *Abstraction by omitting neurons before the l -th layer.*:

The following result states that one can retain soundness for safety verification, by considering all possible neuron values that can appear in the l -th layer.

Lemma 1 (Verification by Layer Abstraction). *If there exists no $\hat{n}_l \in \mathbb{R}^{d_l}$ such that $g^{(L)}(g^{(L-1)} \dots (g^{(l+1)}(\hat{n}_l))$ satisfies ψ_{out} and $h_i^\phi(\hat{n}_l) = 1$, then the neural network is safe under input constraint ϕ and output risk constraint ψ .*

Proof. The lemma holds because for every input $\text{in} \in \mathbb{R}^{d_0}$ of the network, $f^{(l)}(\text{in}) \in \mathbb{R}^{d_l}$. \square

Obviously, the use of \mathbb{R}^{d_l} in Lemma 1 is overly conservative, and we can strengthen Lemma 1 without losing soundness, if we find $S \subseteq \mathbb{R}^{d_l}$ which guarantees that $f^{(l)}(\text{in}) \in S$ for every input $\text{in} \in \mathbb{R}^{d_0}$ of the network. Obtaining such a set S can be achieved by abstract interpretation techniques [6], [21] which perform symbolic reasoning over the neural network in a layer-wise manner.

Lemma 2 (Abstraction via Input Over-approximation). *Let $S \subseteq \mathbb{R}^{d_l}$ guarantee that $f^{(l)}(\text{in}) \in S$ for every input $\text{in} \in \mathbb{R}^{d_0}$ of the network. If there exists no $\hat{n}_l \in S$ such that $g^{(L)}(g^{(L-1)} \dots (g^{(l+1)}(\hat{n}_l))$ satisfies ψ_{out} and $h_i^\phi(\hat{n}_l) = 1$, then the neural network is safe under input constraint ϕ and output risk constraint ψ .*

b) *Assume-guarantee Verification via Monitoring.*: If the computed \mathcal{S} , due to over-approximation, is too coarse to prove safety, one practical alternative is to generate \tilde{S} which only guarantees $f^{(l)}(\text{in}) \in \tilde{S}$ for every input $\text{in} \in \text{In}$ in the training data. In other words, \tilde{S} over-approximates the neuron values computed based on the samples in the training data.

If using \tilde{S} is sufficient to prove safety and if for any input in , checking whether $f^{(l)}(\text{in}) \in \tilde{S}$ can be computed efficiently, one can conditionally accept the proof by designing a runtime monitor which raises a *warning* that the assumption $f^{(l)}(\text{in}) \in \tilde{S}$ used in the proof is violated. Admittedly, \tilde{S} can be an under-approximation over $\{f^{(l)}(\text{in}) \mid \text{in} \in \mathbb{R}^{d_0}\}$, but practically creating an over-approximation only based on the training data is useful and can avoid unstructured input such as noise which is allowed when using \mathbb{R}^{d_0} .

III. TOWARDS STATISTICAL REASONING

The results in Section II are based on two assumptions of *perfection*, namely

- **(perfect training)** the input property characterizer perfectly decides whether property ϕ holds, for each sample in the training data, and

	$\text{in} \in \text{In}_\phi$	$\text{in} \notin \text{In}_\phi$
$h_i^\phi(f^{(l)}(\text{in})) = 1$	α	β
$h_i^\phi(f^{(l)}(\text{in})) = 0$	γ	$1 - \alpha - \beta - \gamma$

TABLE I

PROBABILITY BY CONSIDERING ALL POSSIBLE CASES DUE TO DECISIONS MADE BY THE INPUT CHARACTERIZER (WHETHER $h_i^\phi(f^{(l)}(\text{in})) = 1$) AND THE GROUND TRUTH (WHETHER $\text{IN} \in \text{In}_\phi$).

- **(perfect generalization)** the input property characterizer generalizes its decision (whether property ϕ holds) also perfectly to every data point in the complete input space.

One important question appears when the above two assumptions do not hold, meaning that it is possible for the input property characterizer to make mistakes. By considering all four possibilities in Table I, one realizes that even when a safety proof is established by considering all inputs where $h_i^\phi(f^{(l)}(\text{in})) = 1$, there exists a probability γ where an input in should be analyzed, but in is omitted in the proof process due to $h_i^\phi(f^{(l)}(\text{in}))$ being 0 (i.e., $\text{in} \in \text{In}_\phi$ and $h_i^\phi(f^{(l)}(\text{in})) = 0$). Therefore, one can only establish a statistical guarantee with $(1 - \gamma)$ probability over the correctness claim³, provided that all data points used in training h_i^ϕ are also safe⁴.

IV. RELATED WORK

Formal verification of neural networks has drawn huge attention with many results available [13], [8], [3], [5], [9], [11], [6], [4], [1], [14], [20], [19], [21]. Although specifications used in formal verification of neural networks are discussed in recent reviews [7], [15], the specification problem in terms of characterising an image set is not addressed, so research results largely use inherent properties of a neural network such as local robustness (as output invariance) or output ranges where one does not need to characterize properties over a set of input images. Not being able to properly characterizing input conditions (one possibility is to simply consider every input to be bounded by $[-1, 1]$) makes it difficult for formal static analysis to achieve any useful results on deep perception networks, regardless of the type of abstraction domain being used (box, octagon, or zonotope). Lastly, our work is motivated by *zero shot learning* [12] which trains additional features apart from a standard neural network. The feature detector is commonly created by extending the network from close-to-output layers.

V. EVALUATION AND CONCLUDING REMARKS

We have applied this methodology to examine a direct perception neural network developed by Audi. The network acts as a hot standby system and computes the next waypoint

³Note that for parts where $\text{in} \notin \text{In}_\phi$ and $h_i^\phi(f^{(l)}(\text{in})) = 0$, no problem occurs as the safety analysis guarantees the desired property when $h_i^\phi(f^{(l)}(\text{in})) = 1$.

⁴In other words, for every $(\text{in}, c) \in (\text{In}, \mathcal{C}_\phi)$, if $h_i^\phi(f^{(l)}(\text{in})) = 0$ and $c = 1$, then $f^{(L)}(\text{in})$ does not satisfy ψ .

and orientation for autonomous vehicles to follow. As the close-to-output layers of the network are either ReLU or Batch Normalization, and as ψ is a conjunction of linear constraints over output, it is feasible to use exact verification methods such as ReLUplex [8], Planet [5] or MILP-based approaches [3], [9] as the underlying verification method. We developed a variation of nn-dependability-kit⁵ to read models from TensorFlow⁶ and to perform formal verification via a reduction to MILP. Using assume-guarantee based techniques that take an over-approximation from neuron values produced by the training data⁷, it is possible to conditionally prove some properties such as “impossibility to suggest steering to the far left, when the road image is bending to the right”. However, under the current setup, it is still impossible to prove intriguing properties such as “impossibility to suggest steering straight, when the road image is bending to the right”. We suspect that the main reason is due to the inherent limitation of the neural network under analysis.

In our experiment, we also found that for some input properties such as traffic participants in adjacent lanes, it is very difficult to construct the corresponding input property characterizers by taking neuron values from close-to-output layers (i.e., the trained classifier almost acts like fair coin flipping). Based on the theory of *information bottleneck* for neural networks [18], [16], a neural network from high dimensional input to low dimensional output naturally eliminates unrelated information in close-to-output layers. Therefore, the input property can be unrelated to the output of the network. Although we are unable to prove that the output of the network is safe under these input constraints, it should be possible to construct a counter example either by capturing more data or by using adversarial perturbation techniques [17], [10].

To achieve meaningful formal verification, in our experiments, we also realized that it is commonly not sufficient to only record the minimum and maximum value for each neuron, as boxed abstraction can lead to huge over-approximation. In certain circumstances, we also record the minimum and maximum difference between two adjacent neurons in a layer (in Figure 1, we record $n_{i+1}^{17} - n_i^{17}$ where $i \in \{1, 2, 3, 4\}$). Modern training frameworks such as TensorFlow support computing differences of adjacent neurons with GPU parallelization⁸, thereby making monitoring possible.

Overall, our initial result demonstrates the potential of using formal methods even on very complex neural networks, while it provides a clear path to engineers to resolve the problem related to how to characterize input conditions for verification (by also applying machine learning techniques). Our approach of looking at close-to-output layers can be viewed as an abstraction which can, in future work, leads to layer-wise

incremental abstraction-refinement techniques. Although our practical motivation is to verify direct perception networks, the presented technique is equally applicable to any deep network for vision and lidar systems where input constraints are hard to characterize. It opens a new research direction of using learning to assist practical verification of learning systems.

Acknowledgement The research work is conducted during the first author’s service at the fortiss research institute and is supported by the following projects: “Audi Verifiable AI” from Audi AG, Germany and “Dependable AI for automotive systems” from DENSO Corporation, Japan.

REFERENCES

- [1] R. R. Bunel, I. Turkaslan, P. Torr, P. Kohli, and P. K. Mudigonda. A unified view of piecewise linear neural network verification. In *NIPS*, pages 4795–4804, 2018.
- [2] C. Chen, A. Seff, A. Kornhauser, and J. Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *ICCV*, pages 2722–2730. IEEE, 2015.
- [3] C.-H. Cheng, G. Nührenberg, and H. Ruess. Maximum resilience of artificial neural networks. In *ATVA*, pages 251–268. Springer, 2017.
- [4] S. Dutta, S. Jha, S. Sankaranarayanan, and A. Tiwari. Output range analysis for deep feedforward neural networks. In *NFM*, pages 121–138. Springer, 2018.
- [5] R. Ehlers. Formal verification of piece-wise linear feed-forward neural networks. In *ATVA*, pages 269–286. Springer, 2017.
- [6] 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 *Oakland*, pages 3–18. IEEE, 2018.
- [7] X. Huang, D. Kroening, M. Kwiatkowska, W. Ruan, Y. Sun, E. Thamo, M. Wu, and X. Yi. Safety and trustworthiness of deep neural networks: A survey. *arXiv:1812.08342*, 2018.
- [8] G. Katz, C. W. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *CAV*, pages 97–117, 2017.
- [9] A. Lomuscio and L. Maganti. An approach to reachability analysis for feed-forward ReLU neural networks. *arXiv:1706.07351*, 2017.
- [10] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard. Universal adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1765–1773, 2017.
- [11] N. Narodytska, S. Kasiviswanathan, L. Ryzhyk, M. Sagiv, and T. Walsh. Verifying properties of binarized deep neural networks. In *AAAI*, pages 6615–6624, 2018.
- [12] M. Palatucci, D. Pomerleau, G. E. Hinton, and T. M. Mitchell. Zero-shot learning with semantic output codes. In *NIPS*, pages 1410–1418, 2009.
- [13] L. Pulina and A. Tacchella. An abstraction-refinement approach to verification of artificial neural networks. In *CAV*, pages 243–257. Springer, 2010.
- [14] W. Ruan, X. Huang, and M. Kwiatkowska. Reachability analysis of deep neural networks with provable guarantees. In *IJCAI*, pages 2651–2659, 2018.
- [15] S. A. Seshia, A. Desai, T. Dreossi, D. J. Fremont, S. Ghosh, E. Kim, S. Shivakumar, M. Vazquez-Chanlatte, and X. Yue. Formal specification for deep neural networks. In *ATVA*, pages 20–34. Springer, 2018.
- [16] R. Shwartz-Ziv and N. Tishby. Opening the black box of deep neural networks via information. *arXiv:1703.00810*, 2017.
- [17] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv:1312.6199*, 2013.
- [18] N. Tishby and N. Zaslavsky. Deep learning and the information bottleneck principle. In *ITW*, pages 1–5. IEEE, 2015.
- [19] S. Wang, K. Pei, J. Whitehouse, J. Yang, and S. Jana. Formal security analysis of neural networks using symbolic intervals. In *USENIX*, pages 1599–1614, 2018.
- [20] T.-W. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, D. Boning, I. S. Dhillon, and L. Daniel. Towards fast computation of certified robustness for ReLU networks. In *ICML*, pages 5276–5285, 2018.
- [21] P. Yang, J. Liu, J. Li, L. Chen, and X. Huang. Analyzing deep neural networks with symbolic propagation: Towards higher precision and faster verification. *arXiv:1902.09866*, 2019.

⁵<https://github.com/dependable-ai/nn-dependability-kit>

⁶<https://www.tensorflow.org/>

⁷The data is taken from a particular segment of the German A9 highway, by considering variations such as weather and the current lane.

⁸Computing the neuron difference, when neuron values are stored in an 1D tensor n can be done in numpy using a single instruction `diff(n)`, and in TensorFlow using `n[1:] - n[:-1]`.