






# Syntactic vs Semantic Linear Abstraction and Refinement of Neural Networks

Calvin Chau<sup>1</sup> , Jan Křetínský<sup>2,3</sup> , and Stefanie Mohr<sup>2</sup> 

<sup>1</sup> Technische Universität Dresden, Dresden, Germany

`calvin.chau@tu-dresden.de`

<sup>2</sup> Technical University of Munich, Munich, Germany

`{kretinsky,mohr}@in.tum.de`

<sup>3</sup> Masaryk University, Brno, Czech Republic

**Abstract.** Abstraction is a key verification technique to improve scalability. However, its use for neural networks is so far extremely limited. Previous approaches for abstracting classification networks replace several neurons with one of them that is similar enough. We can classify the similarity as defined either syntactically (using quantities on the connections between neurons) or semantically (on the activation values of neurons for various inputs). Unfortunately, the previous approaches only achieve moderate reductions, when implemented at all. In this work, we provide a more flexible framework, where a neuron can be replaced with a *linear combination* of other neurons, improving the reduction. We apply this approach both on syntactic and semantic abstractions, and implement and evaluate them experimentally. Further, we introduce a refinement method for our abstractions, allowing for finding a better balance between reduction and precision.

**Keywords:** Neural network · Abstraction · Machine learning

## 1 Introduction

**Neural Network Abstractions.** Abstraction is a key instrument for understanding complex systems and analyzing complex problems across all disciplines, including computer science. Abstraction of complex systems, such as neural networks (NN), results in smaller systems, which are not only producing equivalent outputs (such as in distillation [13]), but additionally can be mapped to the original system, providing a strong link between the individual parts of the two systems. Consequently, abstraction find various applications. For instance, the smaller (abstract) networks are more understandable and the strong link between the behaviours of the abstract and the original network allows for better explainability of the original behaviour, too; smaller networks are more efficient

---

This research was funded in part by the German Research Foundation (DFG) project 427755713 *GoPro*, the German Federal Ministry of Education and Research (BMBF) within the project *SEMECO Q1* (03ZU1210AG), and the DFG research training group *ConVeY* (GRK 2428).

in resource usage during runtime; smaller networks are easier to verify. Again, with no formal link between the original network and, say, a distilled or pruned one, verifying the smaller one is of no use to verifying the original one. In contrast, for abstractions, the verification guarantee can be in principle transferred to the original network, be it via lifting a counterexample or a proof of correctness.

Altogether, abstractions of neural networks are a key concept worth investigating *eo ipso*, subsequently offering various applications. However, currently it is still very under-developed. For defining an abstraction, we need a transformation linking the original neurons to those in the abstraction. Equivalently, we need a notion of the *similarity of neurons*, to identify a good representative of a group of neurons. The difficulty in contrast to, e.g., predicate abstraction of programs is that neurons have no inner structure such as values of variables stored in a state. On the one hand, approaches based on bisimilarity [22] offer a solution focusing on the “*syntax*” of neurons: the weights of the incoming connections. The quantities give rise to an equivalence akin to probabilistic bisimulation. On the other hand, in search of a stronger tool, approaches such as [2] try to identify “*semantics*” of the neurons. For instance, given a vector of inputs to the network, the *I/O semantics* of a neuron [2] is the vector of activation values of this neuron obtained on these inputs. This represents a finite-dimensional approximation of the actual semantics of a neuron as a computational device. Either way, replacing several neurons with one that is very similar yields only moderate savings on size if the abstract network is supposed to be similar, i.e., yield mostly the same predictions and ensure a tight connection between the similar neurons.

**Our Contribution.** We focus on studying abstraction irrespective of the use case (verification, smaller networks, explainability), to establish a better principal understanding of this crucial, yet in this context underdeveloped technique. First, we explore a richer abstraction scheme, where a group of neurons can be represented not only by a chosen neuron but also by a *linear combination of neurons*. Thus instead of keeping exactly one representative per group, we can “reuse” the chosen representatives in many linear combinations; in other words, the representatives can attain many roles, partially representing many groups, which reduces their required count. We provide several algorithms to do so, ranging from resource-intensive algorithms aiming to show the limits of the approach to efficient heuristics approximating the former ones quite closely. We apply these algorithms to the semantic approach of [2] as well as to the syntactic, bisimulation-like approach similar to [22] not implemented previously. Experimental results confirm the *greater power of this linear-combination* approach; further, they provide insight into the *advantages of semantic similarity over the syntactic* one, pointing out the more advantageous future research directions.

Further, we provide a formal link between the concrete and abstract neurons by proving an error bound induced by the abstraction, showing the abstraction is valid and (approximately) simulates the original network. We show the bound is better than the one based on bisimulation. While still not very practical, the experiments show that even on unseen data, the error is always closely bounded by the error on the data used for generating the abstraction, and mostly even

a lot smaller. This empirical version of the concept of error could thus enable the transfer of reasoning about the abstraction to the original network in a yet much tighter way.

In addition, we suggest *abstraction-refinement* procedures to better fine-tune the trade-off between the precision and the size of the abstraction. The experiments reveal that a more aggressive abstraction followed by a refinement provides better results than a direct, moderate abstraction. Hence involving our refinement in the abstraction process improves the resulting quality, opening new lines of attack on efficient neural network abstractions.

**Summary.** Our contribution can be summarized as follows:

- We define abstractions of neural networks with (approximate) equivalences being linear equations over semantics of neurons. We provide a theoretical bound on the induced error, see Theorem 1. We reflect this idea also on the syntactic, bisimulation-based abstraction.
- We implement both approaches and compare them mutually as well as to their previous, special cases with equivalences being (approximate) identities. We perform the experiments on a number of standard benchmarks, such as MNIST, CIFAR, or FashionMNIST, concluding advantages of semantic over syntactic approaches and of linear over identity-based ones.
- We introduce an abstraction-refinement procedure and also evaluate its benefits experimentally.

**Related Work.** There are various approaches for verification of NN, however, we are **not** presenting another verifier. Instead, we introduce an approach that is **orthogonal to verification** and could be integrated with an existing verifier. Therefore, we do not compare our approach to any verification tool and refer the interested reader to the Verification of Neural Networks Competition [4] for an overview of existing approaches [16, 26, 31, 33].

Network compression techniques share many similarities with abstraction [7] and either focus on reducing the memory footprint [14, 15] or computation time of the model [12], but in contrast, do not provide any formal relation to the original network, rendering them inappropriate for understanding redundancies or verification. Knowledge distillation is a prominent technique, which can reduce networks by a significant amount, but completely loses any connection to the original network [13], and can thus not be used in verification. There is some progress in using abstract domains for scalable verification, like [26, 27, 29], but they do not produce an abstracted NN for verification. Instead, they apply abstraction only tightly entangled together with the verification algorithm. These approaches also try to generate a more scalable verification, however, the key difference is that they do not return an actual abstracted network that could be reused or manually inspected. Katz et al. [8] introduce an abstraction scheme for NN, in which they decompose neurons into several parts, before merging them again to obtain an over-approximation of the original network. However, their approach is limited to networks with one output neuron. For networks with more output neurons, the property to be verified needs to be baked into the

network, making the approach significantly less flexible. Additionally, this tight entanglement of specification and neural network does not allow for retrieving the abstraction later and reusing it for anything else than to verify that specific property. This strongly contrasts our generic and usage-agnostic abstraction and their property-restricted abstractions.

Some other works use abstraction after representing a neural network as an interval neural network [23], or more generally, by using more complex abstract domains [28]. While theoretically interesting, the practicality of these works has not been investigated. There are two approaches that we consider to be the closest to our work: a bisimulation-based approach [22], and *DeepAbstract* [2], which we will more closely introduce in the preliminaries, and compare to in the experiments.

## 2 Preliminaries

In this work, we focus on classification feedforward neural networks. Such a neural network  $N$  consists of several layers  $1, 2, \dots, L$ , with 1 being the *input layer*,  $L$  being the *output layer* and  $2, \dots, L - 1$  being the *hidden layers*. Each layer  $\ell$  contains  $n_\ell$  neurons. Neurons of one layer are connected to neurons of the previous and next layers by means of weighted connections. Associated with every layer  $\ell$  that is not an output layer is a *weight matrix*  $W^{(\ell)} = (w^{(\ell)}(i, j)) \in \mathbb{R}^{n_{\ell+1} \times n_\ell}$  where  $w^{(\ell)}(i, j)$  gives the weights of the connections to the  $i^{\text{th}}$  neuron in layer  $\ell + 1$  from the  $j^{\text{th}}$  neuron in layer  $\ell$ . We use the notation  $W_{i,*}^{(\ell)} = [w^{(\ell)}(i, 1), \dots, w^{(\ell)}(i, n_\ell)]$  to denote the incoming weights of neuron  $i$  in layer  $\ell + 1$  and  $W_{*,j}^{(\ell)} = [w^{(\ell)}(1, j), \dots, w^{(\ell)}(n_{\ell+1}, j)]^\top$  to denote the outgoing weights of neuron  $j$  in layer  $\ell$ . Note that  $W_{i,*}^{(\ell)}$  and  $W_{*,j}^{(\ell)}$  correspond to the  $i^{\text{th}}$  row and  $j^{\text{th}}$  column of  $W^{(\ell)}$  respectively. A vector  $\mathbf{b}^{(\ell)} = [b_1^{(\ell)}, \dots, b_{n_\ell}^{(\ell)}] \in \mathbb{R}^{n_\ell}$  called *bias* is also associated with each hidden layer  $\ell$ . The input and output of a neuron  $i$  in layer  $\ell$  is denoted by  $h_i^{(\ell)}$  and  $z_i^{(\ell)}$  respectively. We call  $\mathbf{h}^\ell = [h_1^{(\ell)}, \dots, h_{n_\ell}^{(\ell)}]^\top$  the vector of *pre-activations* and  $\mathbf{z}^\ell = [z_1^{(\ell)}, \dots, z_{n_\ell}^{(\ell)}]^\top$  the vector of *activations* of layer  $\ell$ . The neuron takes the input  $\mathbf{h}^\ell$ , and applies an *activation function*  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  element-wise on it. The output is then calculated as  $\mathbf{z}^\ell = \phi(\mathbf{h}^\ell)$ , where standard activation functions include tanh, sigmoid, or ReLU [21]. We assume that the activation function is Lipschitz continuous, which in particular holds for the aforementioned functions [30]. In a feedforward neural network, information flows strictly in one direction: from layer  $\ell_m$  to layer  $\ell_n$  where  $\ell_m < \ell_n$ . For an  $n_1$ -dimensional input  $\mathbf{x} \in \mathcal{X}$  from some input space  $\mathcal{X} \subseteq \mathbb{R}^{n_1}$ , the output  $\mathbf{y} \in \mathbb{R}^{n_L}$  of the neural network  $N$ , also written as  $\mathbf{y} = N(\mathbf{x})$  is iteratively computed as:

$$\begin{aligned} \mathbf{h}^{(0)} &= \mathbf{z}^{(0)} = \mathbf{x} \\ \mathbf{h}^{(\ell+1)} &= W^{(\ell)} \mathbf{z}^{(\ell)} + \mathbf{b}^{(\ell+1)} \end{aligned} \tag{1}$$

$$\mathbf{z}^{(\ell+1)} = \phi(\mathbf{h}^{(\ell+1)}) \tag{2}$$

$$\mathbf{y} = \mathbf{z}^{(L)}$$

where  $\phi(\mathbf{x})$  is the column vector obtained by applying  $\phi$  component-wise to  $\mathbf{x}$ . We abuse the notation and write  $\mathbf{z}^{(\ell)}(\mathbf{x})$ , when we want to specify that the output of layer  $\ell$  is computed by starting with  $\mathbf{x}$  as input to the network.

## 2.1 Syntactic and Semantic Abstractions

We are interested in a general abstraction scheme that is not only useful for verification, but also for revealing redundancies, while keeping a formal link to the original network. We distinguish between two types of abstraction: semantic and syntactic. Syntactic abstraction makes use of the weights of the network, the syntactic information, and allows for overapproximation guarantees that are not restricted to specific inputs. However, as we shall see in the experiments, the semantic abstraction can capture the behavior of the original network on typical input data much more accurately than its syntactic counterpart. This comes at the cost of a more challenging error analysis.


**Semantic Information.** In line with *DeepAbstract* [2], we will create the semantic information based on a set of inputs, the *I/O set*,  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathcal{X}$ , which is typically a subset of the training dataset. We use the inputs  $\mathbf{x}_j \in X$ , feed them to the network and store the output values  $\{\mathbf{z}^{(\ell)}(\mathbf{x}_j)\}_{\mathbf{x}_j \in X}$  of a layer  $\ell$  in a matrix  $\mathbf{Z}^{(\ell)} = (z_i^{(\ell)}(\mathbf{x}_j))_{i,j}$ . Note that the columns are the  $\mathbf{z}^{(\ell)}(\mathbf{x}_j)$  and the rows, denoted as  $\mathbf{Z}_{i,*}^{(\ell)}$ , correspond to the values one neuron  $i$  produces for all inputs  $\mathbf{x}_j$ . We refer to the vector  $\mathbf{Z}_{j,*}^{(\ell)}$  as the *semantics* of neuron  $i$ . This collection of matrices  $\mathbf{Z}^{(\ell)}$  for all layers contains the semantic information of the network.

**DeepAbstract.** Since we will compare our approach to *DeepAbstract* [2], we will give a concise description of the idea of their work. First, it generates the semantic information  $\mathbf{Z}$ . For one layer  $\ell$ , it clusters the rows of the matrix by using standard clustering techniques, e.g. k-means clustering [3]. Each cluster is considered to be a group of neurons that have similar semantics and similar behavior. Thus, only one group representative is chosen to remain and the rest is replaced by the representatives.

**Bisimulation.** The idea of [22] is to apply the notion of bisimulation to NN. A bisimulation declares two neurons as equivalent if they agree on their incoming weights, biases, and activation functions. Additionally, the paper introduces a  $\delta$ -bisimulation that allows neurons to be equivalent only up to  $\delta$ , i.e. two neurons  $i, j$  of layer  $\ell$  with the same activation function are considered to be  $\delta$ -bisimilar, if for all  $k : |w^{(\ell-1)}(i, k) - w^{(\ell-1)}(j, k)| \leq \delta$  and  $|b_i^{(\ell)} - b_j^{(\ell)}| \leq \delta$ .

## 3 Linear Abstraction

Our abstraction of a NN is based on the idea that huge NN in their practical application are usually trained with more neurons than necessary. Since there



**Fig. 1.** Linear Abstraction - On the left, the original network with the basis  $B$  in blue. On the right, the abstracted network with the removed neuron  $n_1^1$  and the changed output weights of the basis neurons  $n_2^1, n_3^1$ , where we assume that  $n_1^1$  can be simulated by  $\alpha_{1,1}^{(1)} \cdot n_2^1 + \alpha_{1,2}^{(1)} \cdot n_3^1$ . (Color figure online)

are techniques to avoid “overfitting”, users of machine learning tend to use NN that are bigger than necessary for their task [19]. Intuitively, such networks thus contain redundancies. We want to remove these redundancies to decrease the size of the network and make it more scalable for verification.

Existing approaches group together similar neurons, and then choose a representative. Instead, we propose to replace a neuron with a linear combination of other neurons. More specifically, we want to replace a neuron  $i$  of layer  $\ell$ , not by one single neuron  $j$ , but rather by a clever combination of several neurons, called the *basis*,  $B^{(\ell)} \subset \{1, \dots, n_\ell\} \setminus \{i\}$ , which is a subset of all neurons of this layer and in this case given as their indices. We assume that the behavior of a neuron can be simulated by a linear combination of the behavior of the basis neurons, i.e. by  $\sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} \cdot \mathbf{z}_{j,*}^{(\ell)}$  for some  $\alpha_{i,j}^{(\ell)} \in \mathbb{R}$ .

**Example.** Consider the neural network in Fig. 1. It has an input layer with two neurons  $n_1^0, n_2^0$ , one hidden layer with three neurons  $n_1^1, n_2^1, n_3^1$ , and an output layer with two neurons  $n_1^2, n_2^2$ . We assume that we are given the basis  $B^{(1)} = \{n_2^1, n_3^1\}$ , marked with blue color in the figure, and the linear coefficients  $\alpha_{1,1}^{(1)}, \alpha_{1,2}^{(1)}$ . That is, we assume that  $n_1^1$  can be simulated by the linear combination  $\alpha_{1,1}^{(1)} \cdot n_2^1 + \alpha_{1,2}^{(1)} \cdot n_3^1$ . We can remove neuron  $n_1^1$  and its outgoing weights  $[1, 2]^\top$ , and add the outgoing weights scaled by the linear coefficients to the basis neurons instead. We add  $\alpha_{1,1}^{(1)} \cdot [1, 2]^\top$  to the outgoing weights of neuron  $n_2^1$ , so we get  $[-1, 3]^\top + \alpha_{1,1}^{(1)} \cdot [1, 2]^\top = [-1 + \alpha_{1,1}^{(1)} \cdot 1, 3 + \alpha_{1,1}^{(1)} \cdot 2]^\top$ , and respectively, we get  $[-2 + \alpha_{1,2}^{(1)} \cdot 1, 1 + \alpha_{1,2}^{(1)} \cdot 2]^\top$  as the outgoing weights of neuron  $n_3^1$ .

The computational overhead to compute a linear combination compared to finding a representative is negligible, as we will see in our experiments (see Sect. 5.2). On the other hand, they provide more expressive power, subsuming the aforementioned clustering-based approach [2]. In particular, we can detect scaled weights that previous approaches failed to identify.

Please note that although it is possible to replace a neuron with a linear combination of any other neurons in the network, we will only use neurons

from the same layer due to more efficient support by modern neural network frameworks.

In the following sections, we will answer three questions: How can one find a set of neurons that serves as a basis (Sect. 3.1)? How to find the coefficients for the linear combination (Sect. 3.2)? How to replace a neuron, once its representation as a linear combination is given (Sect. 3.3)?

### 3.1 Finding the Basis

Our approach is meant to find a sufficient smaller subset of neurons in one layer, which is enough to represent the behavior of the whole layer. We will make use of the semantic information of a layer  $\ell$ , given as  $\mathbf{Z}^{(\ell)} = (z_i^{(\ell)}(\mathbf{x}_j))_{i,j}$  (see Sect. 2.1). Based on this, we try to find a basis of neurons, i.e. a set of indices for neurons in this layer  $\{j_1, \dots, j_{k_\ell}\} = B^{(\ell)} \subset \{1, \dots, n_\ell\}$ , which can represent the whole space as well as possible. To this end we want to find a subset of size  $k = |B^{(\ell)}|$  such that  $\|\sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} \cdot \mathbf{Z}_{j,*}^{(\ell)} - \mathbf{Z}_{i,*}^{(\ell)}\|$  is minimized. We denote with

$$A_B = \begin{bmatrix} | & & | \\ \mathbf{Z}_{j_1,*}^{(\ell)} & \dots & \mathbf{Z}_{j_{k_\ell},*}^{(\ell)} \\ | & & | \end{bmatrix} \quad (3)$$

the matrix containing the activations  $\mathbf{Z}_{j,*}^{(\ell)}$  of the neurons in the basis as columns.

**Greedy Algorithm.** The problem of finding an optimal basis of size  $k$  w.r.t.  $L_2$  distance can be seen as a variation of the *column subset selection problem* which is known to be NP-complete [25]. As a consequence, we use a variant of a greedy algorithm [1]. While it does not always yield the optimal solution, it has been observed to work reasonably well in practice [9, 10].

It has already been observed that layers closer to the output usually contain more condensed information and more redundancies, and can, thus, be compressed more aggressively [2]. We present a greedy algorithm that chooses which layer contains more information and needs a larger basis instead of decreasing the basis sizes equally fast in each layer.

In Algorithm 1, we see that the procedure iteratively removes neurons from the basis. To this end, it iterates over all layers  $l \in \{1, \dots, L\}$  in the network. It tries to remove one neuron at a time from the basis. Then it computes the projection error of the smaller basis, which is defined as  $\|\mathbf{Z}^{(\ell)\top} - \Pi_{A_B} \mathbf{Z}^{(\ell)\top}\|$ , where  $\Pi_{A_B}$  is the matrix that projects the columns of  $\mathbf{Z}^{(\ell)\top}$  onto the column space of  $A_B$ . The columns of  $A_B$  are the rows of  $\mathbf{Z}^{(\ell)}$  whose neurons belong to  $B$ . It greedily evaluates all neurons in all layers and selects the best neuron of the best layer to be removed. After checking every layer, the algorithm decides on the best layer and neuron to be removed, i.e. the one with the smallest error.

Since the approach thoroughly evaluates all possibilities, its runtime depends on both the number of layers and neurons. A natural alternative would be a heuristic that guides us similarly well through the search space. We provide our choice of heuristic below.

**Algorithm 1.** Greedy algorithm over all layers

---

```

1: Given:  $k$  neurons to be removed
2:  $\forall l \in \{1, \dots, L\} : B^{(\ell)} \leftarrow \{1, \dots, n_l\}$ 
3:  $error_{min} \leftarrow \infty, l_{best} \leftarrow -1, n_{best} \leftarrow -1$ 
4: for  $i \in 1, \dots, k$  do
5:   for  $l \in 1, \dots, L$  do
6:     for  $j \in 0, \dots, n_l$  do
7:       Compute the projection error  $error_j$  of  $A_{B^{(\ell)} \setminus \{j\}}$ 
8:       if  $error_j < error_{min}$  then
9:          $l_{best} \leftarrow l$ 
10:         $n_{best} \leftarrow j$ 
11:         $error_{min} \leftarrow error_j$ 
12:    $B^{l_{best}} \leftarrow B^{l_{best}} \setminus \{n_{best}\}$ 
13: return  $B^1, \dots, B^L$ 

```

---

**Variance-Based Heuristic.** Instead of a step-wise decision that takes a lot of computation time, we propose to use a variance-based heuristic. We define the variance of a vector  $\mathbf{v} \in \mathbb{R}^n$  in the usual way by  $\text{Var}(\mathbf{v}) = \sum_{i=0}^n (v_i - \text{Mean}(\mathbf{v}))^2$  where  $\text{Mean}(\mathbf{v})$  is the mean of the vector values. W.l.o.g. let the neurons be numbered in such a way that  $\text{Var}(\mathbf{z}_1^{(\ell)}) \geq \dots \geq \text{Var}(\mathbf{z}_{n_\ell}^{(\ell)})$ . We then choose the basis to contain the neurons with the  $k_\ell$  largest variances, i.e.  $B = \{1, \dots, k\}$ . We assume that neurons with a higher variance in their output values carry more information, and are, therefore, more relevant. Indeed, we can see in our experiments, i.e. Fig. 2, that the heuristic-based approach can achieve similar results, but in far less time.

### 3.2 Finding the Coefficients

Given a basis  $B^{(\ell)}$  for some layer  $\ell$ , computed with the before-mentioned approach, we want to find the coefficients that can be used to replace the remaining neurons which are not part of the basis. We fix a neuron  $i$  in layer  $\ell$  that we want to replace and whose values are stored in  $\mathbf{Z}_{i,*}^{(\ell)}$ , and we want to minimize  $\|\sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} \cdot \mathbf{Z}_{j,*}^{(\ell)} - \mathbf{Z}_{i,*}^{(\ell)}\|$  for  $\alpha_{i,j}^{(\ell)}$ .

Since we want to find a linear combination of vectors, a natural choice is **linear programming**. The linear program is straightforward and can be found in [6, Appendix C]. Note that with the linear program, we are minimizing the  $L_1$ -distance between the neuron's values and its replacement, i.e.  $\|\sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} \cdot \mathbf{Z}_{j,*}^{(\ell)} - \mathbf{Z}_{i,*}^{(\ell)}\|_1$ .

In a different way, we can also consider the vectors  $\mathbf{Z}_{j,*}^{(\ell)}$  for  $j \in B^{(\ell)}$  to span a vector space. If we are given a subset  $\{\mathbf{Z}_{j,*}^{(\ell)} | j \in B^{(\ell)} \subset \{1, \dots, n_\ell\}\}$  that forms a basis for this space, i.e.  $\text{span}((\mathbf{Z}_{j,*}^{(\ell)})_{j \in B^{(\ell)}}) = \text{span}((\mathbf{Z}_{j,*}^{(\ell)})_{j \in \{1, \dots, n_\ell\}})$ , we can represent any other vector  $\mathbf{z}_i^{(\ell)}$  in terms of this basis. However, we usually cannot represent one neuron perfectly by a linear combination of other neurons. **Orthogo-**



**nal projection** gives us the closest point in the subspace  $\text{span}((\mathbf{Z}_{j,*}^{(\ell)})_{j \in B^{(\ell)}})$  for any vector, in terms of  $L_2$ -distance. Then,  $\boldsymbol{\alpha} = [\alpha_{i,j_1}^{(\ell)}, \dots, \alpha_{i,j_{k_\ell}}^{(\ell)}]^\top := (A_B^\top A_B)^{-1} A_B^\top \mathbf{Z}_{i,*}^{(\ell)}$  gives us the coefficients for the orthogonal projection of  $\mathbf{Z}_{i,*}^{(\ell)}$  on the linear space spanned by the columns of  $A_B$ . For a more detailed description of orthogonal projection see e.g. [17, Chapter 6.8]. Note that we assume that the columns of  $A_B$  are linearly independent. If not we can simply replace the respective neurons directly.

### 3.3 Replacement

Assuming, we have a basis  $B^{(\ell)}$  of this layer and we already know the coefficients  $\alpha_{i,j}^{(\ell)} \in \mathbb{R}$  for  $j \in B^{(\ell)}$  that we need to simulate the behavior of neuron  $i$ . This means, we have a linear combination  $\sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} \cdot \mathbf{Z}_{j,*}^{(\ell)}$ , which we want to use instead of neuron  $i$  itself. We will replace the outgoing weights  $W^{(\ell)}$  of this layer, such that for all  $j \in B^{(\ell)}$

$$\tilde{W}_{*,j}^{(\ell)} = [w^{(\ell)}(1, j) + \alpha_{i,j}^{(\ell)} w^{(\ell)}(1, i), \dots, w^{(\ell)}(n_{\ell+1}, j) + \alpha_{i,j}^{(\ell)} w^{(\ell)}(n_{\ell+1}, i)]^\top \quad (4)$$

$$= W_{*,j}^{(\ell)} + \alpha_{i,j}^{(\ell)} W_{*,i}^{(\ell)} \quad (5)$$

Furthermore, we set  $\tilde{W}_{*,i}^{(\ell)} = [0, \dots, 0]^\top$ , and  $\tilde{W}_{i,*}^{(\ell)} = [0, \dots, 0]^\top$ . This means that we will not use the output of neuron  $i$  anymore, but rather a weighted sum of the outputs of neurons in  $B^{(\ell)}$ , and that we will not even compute the value of  $i$ . Additionally, we keep track of the changes we apply to the different neurons with a matrix  $D^{(\ell)} = (d_{j,i}^{(\ell)}) \in \mathbb{R}^{n_\ell \times n_{\ell+1}}$ . Initially,  $D^{(\ell)}$  is 0 and after each replacement, we add  $\alpha_{i,j}^{(\ell)} \cdot w^{(\ell)}(i, i')$  to  $d_{j,i'}^{(\ell)}$  for  $j \in B^{(\ell)}$  and  $i' \in \{1, \dots, n_{\ell+1}\}$ . This is necessary for restoring neurons at a later point.

In the optimal case, the replacement will not change the overall behavior of the neural network. We can derive a the same semantic equivalence from [22] incorporated into our setting:

**Proposition 1 (Semantic Equivalence).** *Let  $N$  be a neural network with  $L$  layers,  $\ell$  a layer of  $N$ ,  $i$  a neuron of this layer, and  $B^{(\ell)} \subset \{1, \dots, n_\ell\} \setminus \{i\}$  a chosen basis. Let  $\tilde{N}$  be the NN after replacing neuron  $i$  by a linear combination of basis vectors with coefficients  $\alpha_{i,j}^{(\ell)}$ , with the procedure as described above.*

*If for all inputs  $\mathbf{x} \in X \subset \mathcal{X}$ ,  $z_i^{(\ell)}(\mathbf{x}) = \sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} z_j^{(\ell)}(\mathbf{x})$ , then  $N$  and  $\tilde{N}$  are semantically equal, i.e. for all inputs  $\mathbf{x} \in X$ ,  $\tilde{N}(\mathbf{x}) = N(\mathbf{x})$ .*

It is easy to see that this proposition is true, for a full proof see [6, Appendix A]. However, the proposition assumes equality of  $z_i^{(\ell)}(\mathbf{x})$  and  $\sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} z_j^{(\ell)}(\mathbf{x})$  for  $\mathbf{x} \in X$ , which virtually never holds for real-world neural networks. Therefore, we want to minimize the difference  $|z_i^{(\ell)}(\mathbf{x}) - \sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} z_j^{(\ell)}(\mathbf{x})|$ , which will not yield a semantically equivalent abstraction, but an abstraction with very similar behavior. We can then **quantify the difference** between the output of the

original network and the abstraction, i.e. the *induced error* with the following Theorem.

**Theorem 1 (Over-approximation Guarantee).** *Let  $N$  be an NN with  $L$  layers. For each layer  $\ell$ , we have a basis of neurons  $B^{(\ell)}$ , and a set of replaced neurons  $I^{(\ell)}$ . Then, let  $\tilde{N}$  be the network after replacing neurons in  $I^{(\ell)}$  as described above.*

*We can over-approximate the error between the output of the original network  $N^L$  and the output of the abstraction  $\tilde{N}^L$  for  $\mathbf{x} \in X \subset \mathcal{X}$  by*

$$\|\tilde{N}^L(\mathbf{x}) - N^L(\mathbf{x})\| \leq b(1 - a^{L-1})/(1 - a)$$

*with  $a = \lambda(\|W\| + \eta)$ ,  $b = \lambda\|W\|\epsilon$ , with  $\lambda^{(\ell)}$  being the Lipschitz-constant of the activation function in layer  $\ell$ ,  $\lambda = \max_{\ell} \lambda^{(\ell)}$ ,  $\|W\| = \max_{\ell} \|W^{(\ell)}\|_1$ ,  $\eta = \max_{\ell} \eta^{(\ell)}$ , and  $\epsilon = \max_{\ell} \epsilon^{(\ell)}$ , assuming that for all layers  $\ell \in \{1, \dots, L\}$  and for all inputs  $\mathbf{x} \in X$ , we have*

- for  $i \in I^{(\ell)}$  :  $|z_i^{(\ell)}(\mathbf{x}) - \sum_{j \in B^{(\ell)}} \alpha_{i,j}^{(\ell)} z_j^{(\ell)}(\mathbf{x})| \leq \epsilon^{(\ell)}$
- $|\sum_{i \in I^{(\ell)}} W_{*,i}^{(\ell)} \sum_{t \in B^{(\ell)}} \alpha_{i,t}^{(\ell)}| \leq \eta^{(\ell)}$

In other words, we can over-approximate the difference in the output of the original and the abstracted network by a value that depends on the weight matrices, the activation function and the tightness of the abstracted neurons to their replacements. The proof can be found in [6, Appendix B]. This Theorem provides us with the **theoretical guarantee** that, given our abstraction, we can provide a valid over-approximation of the output of the original network.

**Comparison to the  $\delta$ -Bisimulation.** Let us recap the error definition from [22]. The difference of the bisimulation and the original network is bounded by  $[(2a)^k - 1]b/(2a - 1)$ , where  $a = \lambda|S|\|W\|$  and  $b = \lambda(|P|L(\mathcal{N})\|x\| + 1)\delta^1$ . In this notation,  $|S|$  is the maximum number of neurons per layer in the whole network,  $|P|$  the maximum number of neurons in the bisimulation (can be understood as the number of neurons in an abstraction),  $L(\mathcal{N})$  is the maximum Lipschitz-constant of all layers, and  $\delta$  is the maximum absolute difference of the bias and sum of the incoming weights.

The drawbacks of that approach are twofold: (i) the error is based on one specific input, and (ii) it makes use of the Lipschitz-constant of the whole network. Calculating the Lipschitz constant of an NN is still part of ongoing research [11] and not a trivial problem. In contrast, we improve on both. Our error calculation generalizes over a set of inputs. Additionally, we use local information, stored in the weight-matrices, to circumvent using the Lipschitz-constant of the NN.

## 4 Refinement

For certain inputs the abstraction might not reflect the behavior of the original network. For these inputs, so-called *counterexamples*, we may want to *refine* the

<sup>1</sup> Please note that this statement is slightly different from the paper  $((2a)^k$  instead of  $(2/a)^k$ ), which we believe to be a typo in the paper.

abstraction, as opposed to starting the abstraction from the original network again. We consider an input to be a counterexample whenever the abstraction assigns it a different label than the original network. However, a counterexample can be any input that does not align with the specifications.

We propose to refine the abstraction by restoring some of the replaced neurons. To do this, we need to know which neurons should be replaced and how. We first briefly mention three heuristics to choose a neuron for restoration. Afterward, we explain how to restore a neuron. Note that the refinement offers more than a “roll-back” of the most recent step of the abstraction since it picks the step-to-be-rolled-back in retrospect reflecting all other steps, leading to a more informed choice. This could in principle be done directly in the abstraction phase, but at an infeasible cost of a huge look-ahead.

**Refinement Heuristics.** We propose three different heuristics: difference-guided, gradient-guided, and look-ahead.

- The *difference-guided* refinement looks at the difference of a neuron in the original and its representation as a linear combination in the abstraction. It replaces the neuron with the largest difference.
- The *gradient-guided* refinement additionally takes the gradient of the NN into account, that is computed as in the training phase of the NN. This takes into account how the whole network would need to change to fix the counterexample.
- The *look-ahead* is the most greedy method and would try out every replaced neuron. It would check how much the network would improve if the neuron was replaced and then chooses the neuron with the highest improvement.

More details on the approaches can be found in [6, Appendix D].

**Restoration of a Neuron.** The restoration principle can be seen as the counterpart of the replacement. Let  $\tilde{N}$  be the network obtained by replacing several neurons in the original network  $N$ , where we want to restore a deleted neuron  $i$  of layer  $\ell$ . To do this, we need not only to get the original neuron back, including its incoming and outgoing weights but also to remove the additional outgoing weights from the basis neurons. Intuitively, the restoration removes the linear combination, ensures that the original outgoing weights for the neuron are used, and adjusts the incoming weights of the neuron. We may have changed layer  $\ell - 1$ , and thus we cannot restore the original incoming weights of neuron  $i$ , but we have to adapt it to changes in the basis  $B^{(\ell-1)}$ . This can be done with the following changes:

$$\begin{aligned}
 & - \forall j \in B^{(\ell)}: \tilde{W}_{*,j}^{(\ell)} = \tilde{\tilde{W}}_{*,j}^{(\ell)} - \alpha_j W_{*,j}^{(\ell)} \\
 & - \tilde{\tilde{W}}_{*,i}^{(\ell)} = W_{*,i}^{(\ell)} \\
 & - \forall j \in B^{(\ell-1)}: \tilde{w}^{(\ell-1)}(i, j) = w^{(\ell-1)}(i, j) + d_{j,i}^{(\ell-1)}
 \end{aligned}$$

Afterward, we subtract  $\alpha_j \cdot w^{(\ell)}(i, i')$  from  $d_{j,i'}^{(\ell)}$  for  $i' \in \{1, \dots, n_{\ell+1}\}$  and  $j \in B^{(\ell)}$ .

## 5 Experimental Results

Our experimental section is divided into several parts: The first one covers how the different methods for finding a basis and the coefficients compare, as described in Sect. 3.2 and Sect. 3.1. The second part shows experiments on our approach in comparison to existing works, namely *DeepAbstract* [2] and our implementation of bisimulation [22] (which was not implemented before). The third part contains the comparison between the abstraction based on syntactic and semantic information. The fourth part describes our experiments on abstraction refinement. Finally, the last part contains experiments on the error induced by our abstraction. Note that supplemental experiments can be found in the Appendix.

Lastly, the work of Katz et al. [8] tightly couples the abstraction with the subsequent particular verification, by integrating the specification as layers into the network. It is, thus, not clear how an abstraction from [8] could be extracted from the tool and reused for another purpose. Additionally, our abstraction would have to be connected with some verification algorithm (DeepPoly, as done by DeepAbstract, or some other) to compare. *Any comparison of the two works would then mostly compare the different verification tools, not really the abstractions.* Although a comparison of different verifiers linked to our LiNNA is an interesting next step into one of the possible applications, it is out of the scope of this paper, which examines the abstraction itself (see Introduction).


**Implementation.** We implemented the approach in our tool *LiNNA* (*Linear Neural Network Abstraction*)<sup>2</sup>. We used networks that were trained on MNIST [20], CIFAR-10 [18], and FashionMNIST [32] for our experiments. In the following, we refer to the corresponding trained networks with “ $L \times n$ ”, where  $L$  denotes the number of hidden layers and  $n$  is the number of neurons in these hidden layers. All experiments were conducted on a computer with Ubuntu 22.04 LTS with 2.6 GHz Intel® Core™ i7 processors, and 32 GB of RAM.

**Performance Measures.** We will compare the approaches mostly on (i) the reduction rate and (ii) the accuracy on a test set. Intuitively, the *reduction rate* describes how much the NN was reduced by abstraction. If an NN  $N$  has in total  $n$  neurons, but after reduction, there are  $m$  neurons left, then the reduction rate is then defined as  $RR(N) = 1 - \frac{m}{n}$ . The *accuracy* of a NN on a test set is defined as the ratio of how many inputs are predicted with the correct label. This is the key performance indicator in machine learning and shows how well a network generalizes to unseen data. In evaluating our abstraction, we follow the same principle since we want to know how well the NN generalizes after abstraction. Note that this test set was not used for training or computing the abstraction.

### 5.1 Abstraction

**Finding the Basis.** We have given two different methods in Sect. 3.1 to find a good basis  $B$ . While the orthogonal projection yields an equally good abstraction

<sup>2</sup> <https://github.com/cxlvinchau/LiNNA>.



**Fig. 2.** *Finding the basis for replacement* - Evaluation on different datasets. The plots contain a comparison of LiNNA while using the greedy variant (solid) and the variance-based heuristic (dashed) for finding a basis with orthogonal projection. Comparison of accuracy (blue) in percent and computation time (red) in seconds. (Color figure online)


compared to linear programming, it outperforms the latter in terms of runtime by magnitudes. Hence, we conducted the rest of the experiments with orthogonal projection. The full comparison between orthogonal projection and linear programming can be found in [6, Fig. 14, Appendix E].

When we compare the greedy and the heuristic-based approach, shown in Fig. 2, we see that the former outperforms the latter in terms of accuracy on MNIST and FashionMNIST. On CIFAR-10, the variance-based approach is slightly better. However, the variance-based approach is always faster than the greedy approach and scales better, as can be seen for all datasets. Unsurprisingly, the greedy approach takes more time for higher reduction rates, because it needs to evaluate many candidates. The variance-based approach just takes the best neurons according to their variance, which has to be calculated only once. Therefore, the calculation is constant in terms of removed neurons.


The plots show one more difference in the behavior: On MNIST and FashionMNIST, we see a quite stable accuracy until a reduction rate of 60%. We cannot see the same behavior on CIFAR-10. We believe this is due to the accuracy and size of the networks. Whereas it is fairly easy to train a feedforward network for MNIST and FashionMNIST on a regular computer, this is more challenging for CIFAR-10. We plan to include more extensive experiments including more involved NN architectures in future work. Finally, our abstraction relies on the assumption that NNs contain a lot of redundant information.

We want to emphasize, that in machine learning, it is common to train a huge network that contains many more neurons than necessary to solve the task [34]. After the introduction of regularization techniques (e.g. [24]), the problem of over-fitting (e.g. [5]) has become often negligible. Therefore, the automatic response to a bad neural network is often to increase its size, either in depth or in width. Our approach can detect these cases and abstract away the redundant information.

**Finding the Coefficients.** We have in total four different approaches to finding the coefficients: greedy or heuristic-based linear programming, and greedy or heuristic-based orthogonal projection. All four have similar accuracies for the



**Fig. 3.** Comparison of LiNNA to related work - LiNNA (greedy and heuristic-based variant), *DeepAbstract* [2], and our implementation of the bisimulation [22] is evaluated in terms of accuracy on the test set for a certain reduction rate. The experiment was conducted on an MNIST  $3 \times 100$  network.




**Fig. 4.** Scalability of LiNNA - Average runtime for 20 different reduction rates on one network. The plot at the top depicts the runtime for MNIST networks with 4 layers, w.r.t. number of neurons. The plot at the bottom shows the runtime for MNIST networks with 100 neurons per layer, w.r.t. number of layers.

same reduction rate, whereas the heuristic ones are mostly just slightly worse than the greedy ones. For a more detailed evaluation, please refer to [6, Appendix G]. The runtimes of the four approaches, however, differ a lot. Take for example an MNIST  $3 \times 100$  network. We assume that the abstraction is performed by starting with the full network and reducing up to a certain reduction rate. Thus, we have runtimes for each of the approaches for each reduction rate. We take the average over all the reductions and get 47 s for the greedy orthogonal projection, 5130 s for the greedy linear programming, 1 s for the heuristic orthogonal projection, and 2 s for the heuristic linear programming. Linear programming takes a lot more time than orthogonal projection, and, as already seen before, the heuristic approaches are much faster than the greedy ones. Please refer to [6, Appendix J] for more experiments on the runtime. Therefore, we propose to use the heuristic approach and the orthogonal projection.

**Scalability.** We evaluate how our approach scales to networks of different sizes. We evaluate (1) how our approach scales with an increasing number of layers, and (2) how it scales with a fixed number of layers but an increasing number of neurons. We show our experiments in Fig. 4. The runtime is the average runtime over 20 different reduction rates on the same network. One can imagine this as averaging the runtimes shown in Fig. 2. We can see that the variance-based approach has almost constant runtime, whereas the runtime of the greedy approach is increasing for both a higher number of layers and neurons.

**Final Assessment.** We have four possibilities on how to abstract an NN: greedy orthogonal projection, greedy linear programming, heuristic-based orthogonal



**Fig. 5.** Evolution of the accuracy on the test set for different reduction rates, for an increasing number of layers, or neurons. We show LiNNA (blue-green) for semantic abstraction, and for syntactic abstraction, bisimulation (red-yellow). The networks were trained on MNIST and have a fixed number of neurons (100) on the left, and a fixed number of layers (4) on the right.

projection, and heuristic-based linear programming. Given that the orthogonal projection outperforms linear programming in terms of accuracy and computation time, we propose to use orthogonal projection. We believe that it is sufficient to use the heuristic-based approach, thereby gaining faster runtimes and only barely sacrificing any accuracy. Whenever we refer to *LiNNA* from now on without any additions, it will be the heuristic-based orthogonal projection.


## 5.2 Comparison to Existing Work

We want to show how our approach compares to existing works, i.e. *DeepAbstract* and the *bisimulation*. Since there is no implementation available for the latter, we implemented it ourselves. Please refer to [6, Appendix F] for the details. The results of the comparison are shown in Fig. 3. It is evident that *DeepAbstract* achieves higher accuracies than the bisimulation, but *LiNNA* outperforms *DeepAbstract* and the bisimulation in terms of accuracy for all reduction rates.


Concerning the runtime, we measure the runtime of each approach for a certain reduction rate, starting from the full network. We find that (in the median) *LiNNA* (greedy) needs 55 s up to 199 s, *LiNNA* (heuristic) 2 s up to 3 s, *DeepAbstract* 187 s up to 2420 s, and the bisimulation 1 s up to 2 s, on MNIST networks of different sizes (starting from  $4 \times 50$  up to  $11 \times 100$ ). The details can be found in [6, Appendix J]. The bisimulation performs best, however just slightly ahead of the heuristic-based *LiNNA*. The greedy *LiNNA*, as well as *DeepAbstract* both have a much higher computation time.

However, in terms of accuracy, greedy *LiNNA* seems to be the best-performing approach, given sufficient time. Due to efficiency, we suggest using heuristic-based *LiNNA*, as it is as fast as the bisimulation, but its accuracy is a lot better and even close to greedy *LiNNA*.

Since we are interested in the general behavior of the abstraction, we want to see how the methods work for varying sizes of networks, but not only in terms of scalability. In Fig. 5, we show the trend for bisimulation and *LiNNA* for an increasing number of layers resp. neurons per layer. On the left, we fix the



**Fig. 6. Syntactic VS. Semantic** - This plot shows the difference between using semantic resp. syntactic information for the abstraction on an MNIST  $5 \times 100$  network. Semantic: LiNNA (semantic) and DeepAbstract. Syntactic: LiNNA (syntactic) and the bisimulation.



**Fig. 7. Refinement** - This plot shows the accuracy of an MNIST  $5 \times 100$  network that was abstracted and refined to a certain reduction rate  $R$ . There is also a plot for an abstraction to the same reduction rate as after the refinement but without refining.

number of neurons per layer to 100 and incrementally increase the number of layers. On the right, we fix the number of layers to four and increase the number of neurons.

We can see that the performance of the networks from the bisimulation varies a lot and gets slightly worse when there are more layers, whereas LiNNA has a very small variation and the performance of the abstractions increases slightly for more layers. Both approaches compute abstractions that perform better the more neurons are in a layer, but LiNNA converges to a much steeper curve at high reduction rates.

For NNs with 400 or more neurons, LiNNA can reduce 80% of the neurons without a significant loss in accuracy, whereas the bisimulation can do the same only for up to a reduction rate of 55%.

### 5.3 Semantic vs Syntactic


In the following, we want to show the differences between semantic and syntactic abstractions. Recall that syntactic abstraction makes use of the weights of the network, the syntactic information, with no consideration of the actual behavior of the NN on the inputs. Semantic abstraction, on the other hand, focuses on the values of the neurons on an input dataset, which also incorporates information about the weights. DeepAbstract and LiNNA, both use semantic information, whereas bisimulation uses syntactic information. We additionally evaluate the performance of LiNNA on syntactic information.

*Which type of information is better for abstraction: semantic or syntactic?*


Note that both DeepAbstract and the bisimulation represent a group of neurons by one single representative, whereas LiNNA makes use of a linear combination.

We summarize our results in Fig. 6. For smaller reduction rates, the bisimulation performs better than LiNNA on syntactic information; for higher reduction rates it is reversed. In general, the approaches based on semantics (DeepAbstract





**Fig. 8.** Comparison of refinement techniques on different architectures for MNIST. The respective networks were abstracted with a reduction rate of 50%. The lines show the variance, the box represents 50% of the data, the line in the box shows the median.



**Fig. 9.** Refinement on different layers - We considered abstractions that were obtained with a 50% reduction rate and fixed 1000 counterexamples. The plots depict the percentage of restored neurons in the layers of the different MNIST networks.

and LiNNA - semantic) outperform the other two approaches w.r.t. accuracy. While abstraction based on syntactic information can provide global guarantees for any input, abstraction based on semantic information relies on the fact that its inputs during abstraction are similar to the ones it will be evaluated on later. However, we see that still the semantic information is more appropriate for preserving accuracy because it combines the knowledge about possible inputs with the knowledge about the weights.

#### 5.4 Refining the Network

We propose refinement of the abstraction in cases where it does not capture all the behavior anymore, instead of restarting the abstraction process. We consider networks that are abstracted up to certain reduction rates, i.e. 20%, 30%, ..., 90%, and use the refinement to regain 10% of the neurons. For example, we reduce the network by 90% and then use refinement to get back to a reduction rate of 80%. We evaluate this refined network on the test dataset and plot its accuracy. Additionally, we show the accuracy of the same NN which is directly reduced to an 80% reduction rate, without refinement. This plot is shown in Fig. 7 for a  $5 \times 100$  network, trained on MNIST.

The gradient and look-ahead refinement have a similar performance. However, the difference-based approach even outperforms the direct reduction itself. This behavior can be explained by the fact that the refinement and the abstraction look at different metrics for removing/restoring neurons. The refinement can focus directly on optimizing for the inputs at hand, whereas the abstraction was generated on the training set. In conclusion, the refinement can even improve the abstraction and it is beneficial to abstract slightly more than required, and refine for the relevant inputs, rather than having a finer abstraction directly.


**Comparison of the Different Approaches.** We collect images that are labeled differently by the abstraction and observe the number of neurons that are restored in order to fix the classification of each image. We ran the experiment on different networks that were abstracted with a 50% reduction rate and considered 1000 counterexamples for each network. The results are summarized in Fig. 8, where we have boxplots for each refinement method on four different network architectures. The look-ahead approach is the most effective technique since it requires the smallest number of restored neurons. In the median, it only requires 1 to 2 operations. The gradient-based approach performs noticeably worse but outperforms the difference-based approach on all networks. The computation time, however, gives a different perspective: Repairing one counterexample takes on average <1 s for the difference-based approach, 1 s for the gradient-based, but the look-ahead approach takes on average 4 s. Interestingly, the look-ahead approach restores fewer neurons but performs worse in accuracy. The difference-based performs better in terms of accuracy while restoring more neurons.

**Insight on the Relevance of Layers.** We also investigated in which layers the different refinement techniques tend to restore the neurons. The plots in Fig. 9 illustrate the percentage of restored neurons in each layer. Notably, the look-ahead approach restores most neurons in the first layer, and very few or none in the later layers, whereas the other approaches have a more uniform behavior. However, the more layers the network has, the more the gradient- and difference-based approaches tend to restore more neurons in the first layer. As reported already by [2], the first layers seem to have a larger influence on the network’s output and hence should be focused on during refinement. It is even more interesting that the difference-based approach does not focus on the first layers as much as the look-ahead approach, but it is better in terms of accuracy.

## 5.5 Error Calculation

We want to show how the abstraction simulates the original network on unseen data not only w.r.t. the output but on every single neuron. In other words, is the discrepancy between the concrete and abstract network higher on the *test* data than on the *training* data that generate the abstraction, or does the link between the neuron and its linear abstraction generalize well?

In Fig. 10, we look at this ratio (“relative error of the abstraction”), i.e. the absolute difference of (activation values of) a simulated abstract neuron to the original neuron, once on the test dataset divided by the maximum value on the training dataset. We can see that there are cases where the error can be greater than one (meaning “larger than on the training set”), see the first row of the plot. However, the geometric mean, defined as  $(\prod_{i=1}^N a_i)^{\frac{1}{N}}$ , calculated over all images is very small. Note that more experiments can be found in [6, Appendix L]. In conclusion, we can say that our abstraction is close to the original also on the test dataset, although the theoretical error calculation does not guarantee so tight a simulation. Future work should reveal how to further utilize the empirical proximity in transferring the reasoning from the abstraction to the original.



**Fig. 10.** Histograms of the relative error of the values of the neurons in an MNIST  $3 \times 100$  network and its abstraction (reduced by 30%). The first row shows the maximum relative error of each neuron in the NN, that occurred for some input from the test set. The second row shows the geometric mean of the relative error of each neuron over 100 images of the test set.

## 6 Conclusions

The focus of this work was to examine abstraction not as a part of a verification procedure, but rather as a stand-alone transformation, which can later be used in different ways: as a preprocessing step for verification, as means of obtaining an equivalent smaller network, or to gain insights about the network and its training, such as identifying where redundancies arise in trained neural networks. (This is analogous to the situation of bisimulation, which has been largely investigated on its own not necessarily as a part of a verification procedure, and its use in verification is only one of the applications.)

We have introduced *LiNNA*, which abstracts a network by replacing neurons with *linear combinations* of other neurons and also equip it with a *refinement* method. We bound the error and thus the difference between the abstraction and the original network in Theorem 1. The theorem yields a lower and an upper bound on the network’s output, thereby providing its over-approximation.

We showed that the linear extension provides better performance than existing work on abstraction for classification networks, both DeepAbstract, and the bisimulation-based approach. We focused our experimental evaluation on *accuracy*, since the aim of the abstraction is to faithfully mimic the *whole classification process* in the smaller, abstract network, not just one concrete property to be verified, which describes only a very specific aspect of the network. Interestingly, the practical error is dramatically smaller than the worst-case bounds. We hope this first, experimental step will stimulate interest in research that could utilize this actual advantage, which is currently not supported by any respective theory.

Furthermore, we show that the use of *semantic information should be preferred* over syntactic information because it allows for higher reductions while preserving similar behavior and being cheap since the I/O sets can be quite small. Bringing back semantics could take us closer to the efficiency of classical software abstraction, where the semantics of states is the very key, going way beyond bisimulation.

## References

1. Altschuler, J., et al.: Greedy column subset selection: new bounds and distributed algorithms. In: Balcan, M., Weinberger, K.Q. (eds.) Proceedings of the 33rd International Conference on Machine Learning, ICML, New York City, NY, USA, vol. 48, pp. 2539–2548. JMLR Workshop and Conference Proceedings. JMLR.org (2016)
2. Ashok, P., Hashemi, V., Křetínský, J., Mohr, S.: DeepAbstract: neural network abstraction for accelerating verification. In: Hung, D.V., Sokolsky, O. (eds.) ATVA 2020. LNCS, vol. 12302, pp. 92–107. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59152-6\\_5](https://doi.org/10.1007/978-3-030-59152-6_5)
3. Bishop, C.M.: Pattern Recognition and Machine Learning. Information Science and Statistics, 5th edn. Springer, Cham (2007)
4. Brix, C., et al.: First three years of the international verification of neural networks competition (VNN-COMP). Int. J. Softw. Tools Technol. Transfer 1–11 (2023). <https://doi.org/10.1007/s10009-023-00703-4>
5. Caruana, R., Lawrence, S., Giles, C.: Overfitting in neural nets: backpropagation, conjugate gradient, and early stopping. In: Leen, T., Dietterich, T., Tresp, V. (eds.) Advances in Neural Information Processing Systems, vol. 13. MIT Press (2000)
6. Chau, C., Křetínský, J., Mohr, S.: Syntactic vs semantic linear abstraction and refinement of neural networks (2023)
7. Cheng, Y., Wang, D., Zhou, P., Zhang, T.: A survey of model compression and acceleration for deep neural networks. Preprint [arXiv:1710.09282](https://arxiv.org/abs/1710.09282) (2017)
8. Elboher, Y.Y., Gottschlich, J., Katz, G.: An abstraction-based framework for neural network verification. In: Lahiri, S.K., Wang, C. (eds.) CAV 2020. LNCS, vol. 12224, pp. 43–65. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-53288-8\\_3](https://doi.org/10.1007/978-3-030-53288-8_3)
9. Farahat, A.K., Ghodsi, A., Kamel, M.S.: A fast greedy algorithm for generalized column subset selection. Preprint [arXiv:1312.6820](https://arxiv.org/abs/1312.6820) (2013)
10. Farahat, A.K., Ghodsi, A., Kamel, M.S.: An efficient greedy method for unsupervised feature selection. In: 11th International Conference on Data Mining, Vancouver, BC, Canada, pp. 161–170. IEEE (2011)
11. Fazlyab, M., et al.: Efficient and accurate estimation of Lipschitz constants for deep neural networks. In: Wallach, H., et al. (eds.) Advances in Neural Information Processing Systems, vol. 32. Curran Associates Inc. (2019)
12. Gong, Y., Liu, L., Yang, M., Bourdev, L.: Compressing deep convolutional networks using vector quantization. Preprint [arXiv:1412.6115](https://arxiv.org/abs/1412.6115) (2014)
13. Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. In: NeurIPS Deep Learning Workshop (2014)
14. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4700–4708 (2017)
15. Jian, X., Jinyu, L., Yifan, G.: Restructuring of deep neural network acoustic models with singular value decomposition. In: Interspeech, pp. 2365–2369 (2013). <https://doi.org/10.21437/interspeech.2013-552>
16. Katz, G., et al.: The marabou framework for verification and analysis of deep neural networks. In: Dillig, I., Tasiran, S. (eds.) CAV 2019. LNCS, vol. 11561, pp. 443–452. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-25540-4\\_26](https://doi.org/10.1007/978-3-030-25540-4_26)
17. Kirkwood, J.R., Kirkwood, B.H.: Elementary Linear Algebra. Chapman and Hall/CRC (2017)

18. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
19. Lawrence, S., Giles, C., Tsoi, A.: Lessons in neural network training: overfitting may be harder than expected. In: Anon (ed.) Proceedings of the National Conference on Artificial Intelligence, pp. 540–545. AAAI (1997)
20. LeCun, Y.: The MNIST database of handwritten digits (1998). <http://yann.lecun.com/exdb/mnist/>
21. Maas, A.L., Hannun, A.Y., Ng, A.Y.: Rectifier nonlinearities improve neural network acoustic models. [http://robotics.stanford.edu/~amaas/papers/relu\\_hybrid\\_icml2013\\_final.pdf](http://robotics.stanford.edu/~amaas/papers/relu_hybrid_icml2013_final.pdf)
22. Prabhakar, P.: Bisimulations for neural network reduction. In: Finkbeiner, B., Wies, T. (eds.) VMCAI 2022. LNCS, vol. 13182, pp. 285–300. Springer, Cham (2022). [https://doi.org/10.1007/978-3-030-94583-1\\_14](https://doi.org/10.1007/978-3-030-94583-1_14)
23. Prabhakar, P., Rahimi Afzal, Z.: Abstraction based output range analysis for neural networks. In: Wallach, H., et al. (eds.) Advances in Neural Information Processing Systems, vol. 32. Curran Associates Inc. (2019)
24. Schmidhuber, J.: Deep learning in neural networks: an overview. *Neural Netw.* **61**, 85–117 (2015). <https://doi.org/10.1016/j.neunet.2014.09.003>
25. Shitov, Y.: Column subset selection is NP-complete. *Linear Algebra Appl.* **610**, 52–58 (2021). <https://doi.org/10.1016/j.laa.2020.09.015>
26. Singh, G., Gehr, T., Püschel, M., Vechev, M.: An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.* **3**(POPL) (2019). <https://doi.org/10.1145/3290354>
27. Singh, G., Gehr, T., Püschel, M., Vechev, M.T.: Boosting robustness certification of neural networks. In: 7th International Conference on Learning Representations, ICLR, New Orleans, LA, USA. OpenReview.net (2019)
28. Sotoudeh, M., Thakur, A.V.: Abstract neural networks. In: Pichardie, D., Sighireanu, M. (eds.) SAS 2020. LNCS, vol. 12389, pp. 65–88. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-65474-0\\_4](https://doi.org/10.1007/978-3-030-65474-0_4)
29. Tran, H.-D., et al.: Robustness verification of semantic segmentation neural networks using relaxed reachability. In: Silva, A., Leino, K.R.M. (eds.) CAV 2021. LNCS, vol. 12759, pp. 263–286. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-81685-8\\_12](https://doi.org/10.1007/978-3-030-81685-8_12)
30. Virmaux, A., Scaman, K.: Lipschitz regularity of deep neural networks: analysis and efficient estimation. In: Bengio, S., et al. (eds.) Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems, NeurIPS, Montréal, Canada, pp. 3839–3848 (2018)
31. Wang, S., et al.: Beta-CROWN: efficient bound propagation with per-neuron split constraints for neural network robustness verification. In: Ranzato, M., et al. (eds.) Advances in Neural Information Processing Systems, vol. 34, pp. 29909–29921. Curran Associates Inc. (2021)
32. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. Preprint [arXiv:1708.07747](https://arxiv.org/abs/1708.07747) (2017)
33. Xu, K., et al.: Fast and complete: enabling complete neural network verification with rapid and massively parallel incomplete verifiers. In: International Conference on Learning Representations (2021)
34. Zhang, C., et al.: Understanding deep learning requires rethinking generalization. *CoRR*, abs/1611.03530 (2016). <http://arxiv.org/abs/1611.03530>