

DeepSaDe

Learning Neural Networks That Guarantee Domain Constraint Satisfaction

Goyal, Kshitij; Dumancic, Sebastijan; Blockeel, Hendrik

DOI

[10.1609/aaai.v38i11.29109](https://doi.org/10.1609/aaai.v38i11.29109)

Publication date

2024

Document Version

Final published version

Published in

Technical Tracks 14

Citation (APA)

Goyal, K., Dumancic, S., & Blockeel, H. (2024). DeepSaDe: Learning Neural Networks That Guarantee Domain Constraint Satisfaction. In M. Wooldridge, J. Dy, & S. Natarajan (Eds.), *Technical Tracks 14* (11 ed., Vol. 38, pp. 12199-12207). (Proceedings of the AAAI Conference on Artificial Intelligence; Vol. 38, No. 11). Association for the Advancement of Artificial Intelligence (AAAI).
<https://doi.org/10.1609/aaai.v38i11.29109>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

DeepSaDe: Learning Neural Networks That Guarantee Domain Constraint Satisfaction

Kshitij Goyal¹, Sebastijan Dumancic², Hendrik Blockeel¹

¹KU Leuven, Belgium

²Delft University of Technology, The Netherlands

kgoyal40@gmail.com, s.dumancic@tudelft.nl, hendrik.blockeel@kuleuven.be

Abstract

As machine learning models, specifically neural networks, are becoming increasingly popular, there are concerns regarding their trustworthiness, especially in safety-critical applications, e.g., actions of an autonomous vehicle must be *safe*. There are approaches that can train neural networks where such domain requirements are enforced as constraints, but they either cannot guarantee that the constraint will be satisfied by all possible predictions (even on unseen data) or they are limited in the type of constraints that can be enforced. In this work, we present an approach to train neural networks which can enforce a wide variety of constraints and guarantee that the constraint is satisfied by all possible predictions. The approach builds on earlier work where learning linear models is formulated as a constraint satisfaction problem (CSP). To make this idea applicable to neural networks, two crucial new elements are added: constraint propagation over the network layers, and weight updates based on a mix of gradient descent and CSP solving. Evaluation on various machine learning tasks demonstrates that our approach is flexible enough to enforce a wide variety of domain constraints and is able to guarantee them in neural networks.

1 Introduction

Widespread use of state-of-the-art machine learning (ML) techniques has given rise to concerns regarding the trustworthiness of these models, especially in safety-critical and socially-sensitive domains. For example, in autonomous vehicles that employ ML approaches to predict the next action, the actions must be safe. Such domain requirements can often be formulated as logical constraints on combinations of inputs and outputs (e.g., whenever the input satisfies some condition A, the output must satisfy condition B). Crucially, these domain constraints must be satisfied for all possible inputs to the model, not just the training data. This has motivated researchers to develop approaches that can train models that satisfy a given constraint for all possible predictions.

A general approach to enforcing constraints in ML models is to include a regularization term in the cost function, which typically adds a cost for every violation of a constraint in the training set (e.g., Xu et al. (2018); Diligenti, Gori, and Sacca (2017)). Such an approach can reduce the number of

violations in the training set, but it does not necessarily eliminate them. Moreover, even when it does, this does not guarantee that other instances (outside the training set) cannot violate the constraint. Alternatively, for some model types, such as neural networks, the architecture of the model can be chosen in such a way that certain types of constraints are guaranteed to be satisfied for each possible input (not just training data) (Sivaraman et al. 2020; Hoernle et al. 2022). But this is typically possible only for specific combinations of model and constraint types.

This raises the question of whether generality and certainty can be combined. Is it possible to come up with a *generally applicable* approach that guarantees the satisfaction of constraints not only on the training set but on the *entire* input space, and this for any kind of model? A step in this direction was made by Goyal, Dumancic, and Blockeel (2023), who propose a relatively general solution for linear models. Their approach translates the learning problem into a MaxSMT setting. MaxSMT stands for Maximum Satisfiability Modulo Theories. It is an extension of SAT solving that can take background theories into account (e.g., for reasoning about the real numbers) and that distinguishes soft and hard constraints: it returns a solution that satisfies all hard constraints and as many soft constraints as possible. Goyal, Dumancic, and Blockeel (2023) model the requirements as hard constraints and maximize the fit to the data using soft constraints. Their approach works for a wide range of constraint types but only handles linear models, and assumes a bounded input domain.

In this paper, we substantially extend the applicability of that approach by showing how it can be used to train feedforward neural networks. Two key modifications to the network’s architecture and training procedure suffice for achieving this: (1) propagating the constraints over the network layers to the last layer, which involves adding skip connections (He et al. 2016) that copy the input to the penultimate layer and deriving bounds on the penultimate layer from the bounds on the input (Sunaga 1958), and (2) training the network using a hybrid procedure that uses MaxSMT to determine the weights in the output layer and gradient descent for all other weights. We demonstrate that with these changes, neural networks can be trained that have good performance and guarantee the satisfaction of given constraints. In the following, we first describe the problem setting, then

briefly describe the existing approach that we build on, before detailing our approach. Afterward, we compare our approach to related work and evaluate it experimentally.

2 Problem Statement

In this paper, we focus on semantic constraints which constrain the behavior of the model: the predictions are required to adhere to certain requirements, e.g. safety constraints (Katz et al. 2017), structured output constraints (Xu et al. 2018), and hierarchical constraints (Hoernle et al. 2022). Specifically, we focus on domain constraints: constraints that must hold for all instances in the domain. We assume the constraints can be written with universally quantified logic formulas. In particular, we consider the constraints of form:

$$\mathbf{K} : \forall \mathbf{x} \in \times_{i=1}^n [l_i, u_i], \mathbf{x} \models P \Rightarrow f_{\mathbf{w}}(\mathbf{x}) \models C \quad (1)$$

Which states that if an input \mathbf{x} satisfies a condition P , the output must satisfy a condition C . An example of a safety constraint that can be represented in this way is: “if an object comes in front of a moving vehicle, the vehicle must stop”. Solving for such a constraint exactly using specific constraint solvers allows us to enforce the constraint for all possible inputs. The reason for focusing on this type of constraint is mostly practical: the constraint-solving technology we use was found to scale well enough for these types of constraints. It is not a theoretical limitation: any constraint that can be handled effectively by current constraint-solving technology can be handled by the approach we develop. As demonstrated later in section 6, our chosen constraint formulation (equation 1) already provides us with a variety of tasks to work with.

To make the search procedure tractable, and because features in ML problems are typically bounded (e.g., a pixel in an image takes a value in $[0, 255]$), we use bounded domains $\times_{i=1}^n [l_i, u_i]$ instead of \mathbb{R}^n . As a natural choice, the training data $D \subseteq \mathcal{X} \times \mathcal{Y}$ can be used to calculate these bounds, i.e. $l_i = \min_D(\mathbf{x}_i)$ & $u_i = \max_D(\mathbf{x}_i)$. We rely on the SMT solver Z3 (Moura and Bjørner 2008) to solve such constraints. Other prominent constraint-solving paradigms like MILP and CP do not support such constraints over continuous domains (Nethercote et al. 2007).

We are now ready to formulate our problem statement:

Definition 2.1. Learning Problem. Given a training set $D \subseteq \mathcal{X} \times \mathcal{Y}$, a set of domain constraints \mathcal{K} , a loss function \mathcal{L} , and a hypothesis space containing functions $f_{\mathbf{w}} : \mathcal{X} \rightarrow \mathcal{Y}$; find \mathbf{w} such that $f_{\mathbf{w}}$ satisfies constraints in \mathcal{K} and $\mathcal{L}(f_{\mathbf{w}}, D)$ is minimal among all such $f_{\mathbf{w}}$.

$f_{\mathbf{w}}$ is assumed to be a feedforward neural network. The output layer is real-valued without any activation, and a softmax layer is used to calculate class probabilities for classification problems. The language of the constraints in \mathcal{K} is a subset of first-order logic which allows for universal quantifiers. In practice, we focus on the constraints of the form in equation 1, and \mathcal{K} can be a set of multiple such universally quantified constraints.

3 Background - Satisfiability Descent (SaDe)

Satisfiability (SAT) is the problem of finding a solution that satisfies a given Boolean formula, e.g. $\neg a \vee b$. Satisfiability Modulo Theories (SMT) extend SAT such that formulas can be expressed in other theories, such as Real Arithmetic, e.g., $(a + b > 3) \wedge (b < 1)$ with real a and b . Maximum Satisfiability (MaxSMT) generalizes the SMT problem: given a set of hard constraints \mathcal{H} and soft constraints \mathcal{S} , it aims to find a solution that satisfies all constraints in \mathcal{H} and as many as possible in \mathcal{S} .

Our work builds on *SaDe* (Goyal, Dumancic, and Bloock 2023) which is a learning algorithm that can enforce constraints in linear models and guarantee satisfaction. SaDe modifies the parameter update process of the mini-batch gradient descent (Goodfellow et al. 2016). Unlike gradient descent, which updates the solution in the direction that minimizes the loss, SaDe solves a MaxSMT problem to find the solution at each iteration, which is formulated in such a way that its solution satisfies the domain constraint and is close to the solution that gradient descent might lead to. At each iteration, a local search space is defined for the MaxSMT problem around the previous solution. This search space is a fixed-sized n-cube where each edge of the n-cube is a hyperparameter called *maximal step size* that upper bounds the size of the update in each dimension (illustrated for two parameters in figure 1(a)). SaDe iteratively improves the performance while learning solutions that satisfy the constraint. Formulations of the MaxSMT problem and the local search space at each iteration are provided next.

Formulation of MaxSMT Problem: A MaxSMT problem is defined for a batch of instances in each iteration, where the soft constraints encode a certain quality of fit of the model on the instances in the batch, and the domain constraints are hard constraints.

A soft constraint, for a given instance (\mathbf{x}, y) , is defined as a logical constraint that is fulfilled when the prediction of the model $f_{\mathbf{w}}$ for \mathbf{x} is “sufficiently consistent” with the true label y . For regression, given some error e , a soft constraint is defined as: $|y - e| \leq f_{\mathbf{w}}(\mathbf{x}) \leq |y + e|$. For classification, the sign of $f_{\mathbf{w}}(\mathbf{x})$ is assumed to be the indicator of the class and the magnitude indicates the certainty of prediction, the soft constraint takes the form (for a threshold τ):

$$\begin{aligned} f_{\mathbf{w}}(\mathbf{x}) > \tau & \quad \text{if } y = 1 \\ f_{\mathbf{w}}(\mathbf{x}) < -\tau & \quad \text{if } y = -1 \end{aligned}$$

For each instance, multiple soft constraints are formulated for different values of the error e or threshold τ . Satisfying a maximum of these soft constraints, which is what the MaxSMT problem tries to achieve, correlates with minimizing the prediction loss. Thus, the solution to this MaxSMT problem at every iteration reduces the loss while satisfying the domain constraints.

Local Search Space: At every iteration, the MaxSMT problem searches for the next solution in a local search space defined by the n-cube around the previous solution. It is encoded as an additional hard constraint in the MaxSMT problem, with a “box” constraint, which states that the next solution must be inside the axis-parallel box defined by $\hat{\mathbf{w}}$ and




Figure 1: (a.) Illustration of SaDe: each grey quadrant represents the local search space for the MaxSMT problem (*maximal step size*: α), defined by the gradients of the loss, the green points are the solutions found with MaxSMT; (b.) DeepSaDe Architecture: Last layer is updated using the MaxSMT framework and the layers before are updated with gradient descent. The input features relevant to the domain constraint (in green) are mapped to the penultimate layer via skip connections.

$\hat{\mathbf{w}} - \alpha \cdot \text{sgn}(\mathbf{g})$, where the sign function is applied component-wise to a vector (a modified sign function is used where $\text{sgn}(0) = 1$), $\hat{\mathbf{w}}$ is the previous solution, \mathbf{g} is the gradient of the loss at $\hat{\mathbf{w}}$, and α is the maximal step size. The box constraint serves two important purposes. Firstly, it provides a general direction in which the loss is minimized, as the box in each dimension is aligned with the negative gradient. Secondly, it stabilizes learning by limiting the size of the updates. Interestingly, using a regularized loss instead of a standard loss doesn't make a difference, as the constraint remains fulfilled throughout the training steps. With regularization, the regularization factor only becomes active if an instance violates the constraint. However, in SaDe, this never happens during the learning process, making the regularization always 0.

SaDe is limited to training linear models. Training neural networks with the same procedure would require solving the MaxSMT problem with highly non-linear soft constraints, e.g. $|y - e| \leq f_{\mathbf{w}}(\mathbf{x}) \leq |y + e|$ where $f_{\mathbf{w}}(\mathbf{x})$ is the network, which is not possible with state-of-the-art SMT solvers.

4 DeepSaDe: Deep Satisfiability Descent

We now present our approach *DeepSaDe*, which utilizes the MaxSMT framework proposed in Goyal, Dumancic, and Blokkeel (2023) to train neural networks with constraints. DeepSaDe exploits the structure of neural networks, which transform the input domain through a series of non-linear layers before a final linear layer maps it to the output. As the network output only explicitly depends on the last layer, enforcing the constraint on the last layer is sufficient to enforce the constraint on the network. DeepSaDe, therefore, uses batch learning with a hybrid procedure that uses MaxSMT to determine the weights in the last layer and gradient descent for all other weights. In the following, we introduce some notation before we formalize the MaxSMT problem in the last layer, and then detail the learning algorithm.

$f_{\mathbf{w}}$ is a fully-connected neural network with k layers such that $f_{\mathbf{w}}(\mathbf{x}) = h_k(h_{k-1}(\dots h_1(\mathbf{x})\dots))$ for input $\mathbf{x} \in \mathcal{X}$, where h_n is the n^{th} layer. The input to the n^{th} layer is $\mathbf{x}^{(n)} = h_{n-1}(\dots h_1(\mathbf{x})\dots)$, where $\mathbf{x}^{(1)} = \mathbf{x}$ and the latent input space for the h_n is represented by $\mathcal{X}^{(n)}$, where $\mathcal{X}^{(1)} = \mathcal{X}$. The size of a layer, $|h_n|$, is the number of neurons in it. Weight and bias parameters of h_n are denoted by matrices $W^{(n)}$ and $B^{(n)}$ of dimensions $|h_{n-1}| \times |h_n|$ and $1 \times |h_n|$ respectively ($|h_0| = |\mathbf{x}^{(1)}|$). Elements of $W^{(n)}$ and $B^{(n)}$ are referred with lower-case letters.

Formulation of the Maximum Satisfiability Problem at the Last Layer of the Network

To formalize the MaxSMT problem for the last layer, it is necessary to express both the soft constraints and the hard constraint (domain constraint) in terms of the latent inputs $\mathbf{x}^{(k)}$ rather than the original inputs $\mathbf{x}^{(1)}$. The soft constraint at the last linear layer can be formulated in a similar way to what was presented in section 3. Specifically, for a regression task, the constraint is given by $|y - e| \leq h_k(\mathbf{x}^{(k)}) \leq |y + e|$, whereas for classification, it takes the form of:

$$\begin{aligned} h_k(\mathbf{x}^{(k)}) &> \tau & \text{if } y = 1 \\ h_k(\mathbf{x}^{(k)}) &< -\tau & \text{if } y = -1 \end{aligned}$$

To formulate the hard domain constraints, we propose a method for translating each original constraint K (equation 1) into a constraint K' for the MaxSMT problem at the last layer, given the parameters of layers h_1, \dots, h_{k-1} , such that a solution to K' , combined with h_1, \dots, h_{k-1} , is a solution to K .

$$K' : \forall \mathbf{x}' \in \prod_{i=1}^{|h_{k-1}|} [l_i^{(k)}, u_i^{(k)}], \mathbf{x}' \models P \Rightarrow h_k(\mathbf{x}') \models C \quad (2)$$

Here \mathbf{x}' represent the quantified latent variable in the domain of $\mathcal{X}^{(k)}$ bounded by $\prod_{i=1}^{|h_{k-1}|} [l_i^{(k)}, u_i^{(k)}]$. We explain how this translation is achieved in the next two paragraphs. Given a network, with this translation, the last layer of a network can be updated to satisfy the constraint by the network.

Domain Bound Propagation: To formulate K' , we first consider the bounds of the quantified variable for the latent space of $\mathcal{X}^{(k)}$. The bounds of the latent space must be such

that enforcing the constraint within these bounds enforces the constraint on the original input bounds. To construct such latent bounds, we rely on interval arithmetic (Sunaga 1958). This involves calculating the bounds of the output of a layer based on the bounds of the input and the layer parameters, such that any input to the layer takes the output value within the output bounds. Given the lower and upper bounds $l^{(n)}$ and $u^{(n)}$ for the input $\mathbf{x}^{(n)}$ of the layer h_n , following the approach used in Gowal et al. (2018), the bounds for the output $\mathbf{x}^{(n+1)}$ are computed as (more details in Appendix A.1):

$$l_i^{(n+1)} = \text{act}(b_i^{(n)}) + \sum_{j:w_{j,i}^{(n)} \geq 0} w_{j,i}^{(n)} l_j^{(n)} + \sum_{j:w_{j,i}^{(n)} < 0} w_{j,i}^{(n)} u_j^{(n)}$$

$$u_i^{(n+1)} = \text{act}(b_i^{(n)}) + \sum_{j:w_{j,i}^{(n)} \geq 0} w_{j,i}^{(n)} u_j^{(n)} + \sum_{j:w_{j,i}^{(n)} < 0} w_{j,i}^{(n)} l_j^{(n)}$$

Where ‘act’ is the activation function (like ReLu, Sigmoid, and Tanh). Given the bounds of the input space, the bounds for the latent space $\mathcal{X}^{(k)}$ can be calculated recursively. Enforcing a constraint within the latent bounds enforces the constraint on any input within the input bounds.

Identity Mapping of Relevant Features: The translation also takes into account that some domain constraints may be dependent on the input space, e.g., “if an object comes in front of the vehicle, the vehicle must stop”. To encode such conditions in \mathcal{K}' (i.e., $\mathbf{x}' \models P$), we make the relevant features, i.e. the features that are needed to encode the property P , available at the latent space $\mathcal{X}^{(k)}$. For this, we use skip-connections (He et al. 2016) that map these features to the second to last layer h_{k-1} using an identity mapping, as illustrated in figure 1(b). The network, consequently, is no longer fully-connected and these features take the same value as the input. These mapped features become a part of $\mathcal{X}^{(k)}$ and the input property can be expressed at the last layer. These features are identified in advance. Finally, as the output of the network is the same as the output of the last layer, the constraint on the network output can be encoded with $h_k(\mathbf{x}')$. This completes the formulation of \mathcal{K}' .

Learning Algorithm

We now present the algorithm while referring to the pseudocode in algorithm 1. DeepSaDe modifies mini-batch gradient descent (Goodfellow et al. 2016), where forward and backward passes at every iteration (lines 5-7) are kept the same, but the parameter update is split into two parts. First, only the last layer’s parameters are updated to satisfy the domain constraint using the MaxSMT formulation presented earlier (lines 8-23) (details in next paragraph). Second, the earlier layers are updated using gradient descent to optimize predictive loss (lines 27-29). Importantly, after the latter update, the network does not guarantee constraint satisfaction due to changes in the latent space of the last layer, used to formulate the MaxSMT problem. Therefore, the network before the latter update is used for evaluation and a validation set is used to select the best model (lines 24-26).

For updating the last layer, first, a line search is used along the vector that minimizes the prediction loss, and a fixed

Algorithm 1: Deep Satisfiability Descent (DeepSaDe)

input: training data D , validation data V , domain constraints \mathcal{K} , batch size b , epochs e , loss \mathcal{L} , maximal step size α , learning rate η , line search steps s

```

1:  $f_{\hat{W}}(\cdot)$ : Model initialized with a standard approach
2:  $\mathcal{K}' = \text{TranslateDomainConstraint}(\mathcal{K}, f_{\hat{W}}(\cdot))$ 
3: Update  $\hat{W}$  such that  $\hat{W}^{(k)} \models \mathcal{K}'$ 
4: restart = False;  $\hat{W}_{out} = \text{undefined}$ ; Partition  $D$  into
   batches of size  $b$ 
5: for each epoch do
6:   for each batch  $B$  in  $D$  do
7:      $\mathcal{G} = \nabla \mathcal{L}(\hat{W})$ 
8:      $\mathcal{K}' = \text{TranslateDomainConstraint}(\mathcal{K}, f_{\hat{W}}(\cdot))$ 
9:     if restart then
10:      Randomly flip the sign of each element of  $\mathcal{G}^{(k)}$ 
11:     end if
12:     // try line search, if no solution use MaxSMT
13:     let  $W_i = \hat{W}^{(k)} - \eta \mathcal{G}^{(k)} i/s$ , for  $i = 1, 2, \dots, s$ 
14:     if  $\{i \mid W_i \models \mathcal{K}'\} \neq \emptyset$  then
15:        $\hat{W}^{(k)} = \hat{W}_{\max\{i \mid W_i \models \mathcal{K}'\}}$ 
16:     else
17:        $\mathcal{G}_S = \text{sgn}(\mathcal{G}^{(k)})$ 
18:        $\mathcal{S} = \text{SoftConstraints}(h_{k-1}(\dots h_1(B)\dots))$ 
19:        $\mathcal{H} = \{W^{(k)} \in \text{Box}(\hat{W}^{(k)}, \hat{W}^{(k)} - \alpha \mathcal{G}_S)\} \cup \mathcal{K}'$ 
20:       // assuming MaxSMT returns undefined if unsat
21:        $\hat{W}^{(k)} = \text{MaxSMT}(\mathcal{S}, \mathcal{H})$ 
22:     end if
23:     // no solution found, restart in another direction
24:     if  $\hat{W}^{(k)}$  is undefined then
25:       restart = True
26:     else
27:       // remember best solution found
28:       if  $\text{eval}(f_{\hat{W}}, V) > \text{eval}(f_{\hat{W}_{out}}, V)$  then
29:          $\hat{W}_{out} = \hat{W}$ 
30:       end if
31:       // perform gradient descent on earlier layers
32:       for  $n = 1, \dots, k-1$  do
33:          $W^{(n)} = \hat{W}^{(n)} - \eta * \mathcal{G}^{(n)}$ 
34:       end for
35:       restart = False
36:     end if
37:   end for
38: end for
39: return  $f_{\hat{W}_{out}}(\cdot)$ 

```

number of candidate solutions, from furthest to closest, are checked if they satisfy \mathcal{K}' and first one that does is picked (lines 12-14). If this search does not yield a solution, the MaxSMT problem is formulated based on the inputs to the last layer, and a solution is searched in the local search space around the previous solution defined based on the gradients calculated during the backward pass (lines 16-19) (similar to the approach in SaDe). Line search is employed first because checking if a point satisfies a constraint is faster than

searching for the solution in a domain, finding a solution by merely checking a few candidate points speeds up learning.

Sometimes a solution for the last layer cannot be found because the MaxSMT problem could not be solved within the local search space, possibly due to the gradient pointing to a solution space that violates the constraint. In such cases, a restart procedure is initiated where the signs of the gradients are randomly flipped (lines 9-11, 21-23) to randomize the direction of the update. This may slightly decrease predictive performance, but it effectively restarts the learning process when it gets stuck.

Since our approach is iterative, starting from an initial configuration that satisfies the domain constraint is crucial. Otherwise, we may begin in a solution space far from the constrained space, leading to no updates. We ensure this by first initializing the network using a standard method (He et al. 2015) (line 1), and then updating the weights of the last layer that satisfy the translated constraint \mathcal{K}' (lines 2-3). It is important to note that we are solving a satisfiability problem here, not the maximum satisfiability problem, as we are not using any soft constraints for this purpose.

5 Related Work

A standard approach for enforcing constraints in ML models is *regularization*, where a penalty is added to the prediction loss whenever the model makes a prediction that violates the constraint ($(1 - \lambda) * \text{loss} + \lambda * \text{regularization}$). Xu et al. (2018) propose a regularization defined on the weighted model count (WMC) (Chavira and Darwiche 2008) of the constraint defined over the network output. Diligenti, Gori, and Sacca (2017) and Serafini and Garcez (2016) propose a fuzzy logic-based regularization for constraints in first-order logic. There are many other regularization approaches in the literature, e.g. Fischer et al. (2019), Hu et al. (2016), Stewart and Ermon (2017). Regularization can enforce a variety of constraints, but does not guarantee constraint satisfaction. Additionally, high regularization loss with a large value of λ may provide stronger constraint satisfaction but impacts the predictive performance negatively.

Some approaches guarantee constraints by construction but are generally limited to enforcing specific types of constraints. For example, monotonic lattices (Gupta et al. 2016), deep lattice networks (You et al. 2017), and COMET (Sivaraman et al. 2020) are approaches to enforce monotonicity in neural networks; Leino et al. (2022) & Lin et al. (2020) propose approaches to satisfy some safety specifications. A more general approach, MultiplexNet, was proposed in Hoernle et al. (2022). They use a multiplexer layer to satisfy constraints in disjunctive normal form (DNF). However, DNF representation is limiting as certain constraints have worst-case representations in DNF which may lead to exponentially many terms. Additionally, constraints conditioned on the input space cannot be enforced. Another approach, DeepProbLog (Manhaeve et al. 2018), trains neural networks within the ProbLog framework, where constraints can be enforced with ProbLog. However, it is limited to modeling discrete variables, and cannot model regression.

Our work relates to combinatorial optimization approaches as we use a MaxSMT-based approach. These ap-

proaches, however, except Goyal, Dumancic, and Blockeel (2023), are limited to discrete models like decision trees (Günlük et al. 2021; Bertsimas and Dunn 2017; Verwer and Zhang 2019; Demirović et al. 2022) and decision sets (Yu et al. 2021; Ignatiev et al. 2021). Maximum satisfiability, specifically, has also been used in various ML tasks (Berg et al. 2019; Cussens 2012; Malioutov and Meel 2018). None of these, however, focus on training neural networks.

There are approaches that rely on the idea of bound propagation, also used in our work, to train adversarially robust neural networks (Gowal et al. 2018; Zhang et al. 2019). However, the constraints that can be enforced are limited to input-output bounds (if the input is in a given bound, the output should be in a specific bound). Our approach is more general and can, in theory, handle any constraint that can be written as an SMT formula.

Finally, there are approaches to verify if a network satisfies a constraint (Katz et al. 2019; Wang et al. 2018; Bunel et al. 2018). DeepSaDe trains models that do not require verification because constraints are guaranteed by construction.

6 Experiments

We evaluate multiple use cases in various ML tasks with complex domain constraints. We first outline the research questions, then describe the use cases, our evaluation method, and finally the results.

Q1: Can existing methods satisfy constraints for all predictions in practice, even if they don't guarantee it?

Q2: How does the predictive performance of DeepSaDe models compare to the baselines?

Q3: Do the DeepSaDe models have a higher training time compared to the baselines?

Use Cases

The selection of use cases is done in order to demonstrate the flexible representation of constraints in DeepSaDe. The selected use cases tackle a variety of learning problems and handle different types of constraints. We consider constraints in the form presented in equation 1, where P and C are written as SMT formulas over $\{\forall, \exists, \vee, \wedge, \neg, \geq, \leq, =\}$. In principle, any constraint that can be written as an SMT formula and can be solved using existing SMT solvers can be enforced using DeepSaDe. Use cases UC1, UC2 & UC3 are from Goyal, Dumancic, and Blockeel (2023), UC4 is novel, and UC5 is from Xu et al. (2018).

UC1: A multi-target regression to predict 5 household expenses using 13 attributes, with 41417 data instances. We enforce two constraints: “*sum of all the expenses must be smaller than the total household income*” and “*going out expense must be smaller than 5% of the household income*”.

UC2: A binary classification problem of predicting if a person should be given a loan or not based on 13 attributes, with 492 data instances. We enforce the constraint: “*a person with a salary less than 5000\$ and an absent credit history must be denied a loan*”.

UC3: A multiclass classification problem to classify a song to one of 5 music genres based on 13 attributes, with 793 data instances. We enforce the constraint: “*a song by ‘the Beatles’ must be classified as either rock or pop*”.

UC4: A multi-label classification problem of identifying the labels from a sequence of 4 MNIST images. We enforce the constraint: “the sum of the predicted labels must be greater than 10”. 20000 instances are generated by selecting 4 images at random from the MNIST dataset.

UC5: A preference learning problem to predict the preference order of different sushi, with a constraint: “the prediction must have a coherent preference order”. The preference order of 6 out of 10 sushi is used to predict the preference order of the remaining 4 sushi. The dataset contains 4926 instances. The preference ordering over n items is encoded as a flattened binary matrix $\{X_{ij}\}$ where X_{ij} denotes that item i is at position j . Under this encoding, each instance has 36 features and 16 targets.

The use cases are intentionally designed with some training instances that violate the constraints. While ‘fixing’ or ‘removing’ these violations might be possible in certain cases (like UC2), it is not feasible with complex constraints. For instance, in UC1, assume that there is a training instance where the sum of expenses exceeds household income. Correcting this precisely is problematic because it is unclear how much each expense should be reduced. This issue gets tougher when dealing with multiple constraints. Additionally, real-world scenarios might have violations that are hard to even identify, such as using biased data in fairness-sensitive applications. Thus, we evaluate DeepSaDe and the baselines in cases with such violations. More information about the use cases is provided in Appendix A.2. Importantly, we can only impose constraints explicitly based on model parameters: in the example of a self-driving car (section 2), identifying pixels representing a moving object is in itself a predictive task beyond our current scope. Thus, such cases are not included and are left for future research.

Evaluation Methodology

Evaluation: Constraint satisfaction is typically evaluated by the *constraint accuracy* metric used in Xu et al. (2018) and Fischer et al. (2019), which corresponds to the percent of instances where the constraint is not violated by the prediction. Such an evaluation, however, is limited to a finite sample of the population. Hence, as a second measure, we also calculate the *Adversity Index (AdI)* ((Goyal, Dumancic, and Blockeel 2023)), which is the fraction of instances for which a counter-example to the constraint can be constructed in the neighborhood, defined by an l_∞ ball of radius δ around the instance. AdI takes a value between 0 and 1, a higher value of AdI implies that the model violates the constraint on more points similar to data instances. AdI is calculated on full data (training and test) because this provides more instances to evaluate constraint satisfaction; we want the constraints to be satisfied on all data, not only test data. For DeepSaDe, both constraint accuracy and AdI are zero by construction but we still calculate it as a sanity check. We use the neural network verification software *Marabou* ((Katz et al. 2019)) to find counter-examples. However, we only compute the AdI for UC1-3 because *Marabou* only handles inequality constraints and is not suitable for UC4-5.

For predictive performance, we use MSE for UC1 and accuracy for UC2-3. For UC4-5, we use *coherent accuracy*

which is the fraction of instances for which the model predicts the entire configuration correctly, *flattened accuracy* which is the fraction of individually correct binary labels, and the *Jaccard accuracy* which is the average Jaccard index for each multi-label prediction compared to true labels. Performance is only evaluated for instances where the true label does not violate the constraint as there can be data instances in practice where this happens and a comparison with such instances makes the evaluation biased.

Baselines: For UC2 and UC3, we use regularization baselines based on Xu et al. (2018) (SL) and Diligenti, Gori, and Sacca (2017) (SBR). For UC1, we design a custom regularization loss REG (details in appendix A.2). For UC4, we could not find any approaches that can enforce such a constraint. Hence, we simply compare it with a feedforward network (FFN). For UC5, we choose SL and FFN as baselines. In Goyal, Dumancic, and Blockeel (2023), the authors use a post-processing baseline to enforce constraints at inference time. This type of baseline is excluded from our evaluation because, as also noted by Goyal, Dumancic, and Blockeel (2023), this baseline is only applicable when the constraints are simple and not a generally applicable method.

Experimental Setup: For solving the MaxSMT problem, we implement the Fu-Malik algorithm (Fu and Malik 2006) using the Z3 solver for NRA (Quantified Nonlinear Real Arithmetic) formulas. We ran experiments on an Intel(R) Xeon(R) Silver 4214 CPU @ 2.20GHz machine with 125 GB RAM. For each use case, we run 5 experiments with 5-fold cross-validation, and the data is split 70/20/10 into train/test/validation. Every feature is scaled in $[0, 1]$, and the radius $\delta = 0.1$ is chosen for AdI, which is significantly smaller than the mean l_∞ distance between two points: this distance for UC1 is 0.75, for UC2 is 0.97, and for UC3 is 0.89. For regularization, the smallest value of λ , in $[0, 1]$, that leads to minimum violations on the validation set is selected via cross-validation. Refer to appendix A.2 for details on the architectures and hyper-parameters.

Results (Tables 1 & 2)

Constraint Satisfaction: DeepSaDe finds a model that achieves 100% constraint accuracy for each use case and $\text{AdI} = 0$ for UC1-3. For UC1, the constraint accuracy for REG is 93.5%, and counterexamples can be constructed close to 97% of the instances ($\text{AdI} = 0.97$). Similar behavior is seen for UC3 for both SL and SBR, with SBR proving to be more effective in enforcing constraints because of lower AdI. For UC2, SL and SBR both lead to 100% constraint accuracy, but counterexamples can still be constructed as $\text{AdI} > 0$. For UC4 and UC5, FFN fails to satisfy constraints on test set. For UC5, SL can satisfy the constraint but the predictive performance is much worse than DeepSaDe. Thus, in general, the baselines, in contrast to DeepSaDe, do not satisfy domain constraint satisfaction, which answers Q1.

Predictive Performance: DeepSaDe treats the domain constraint as a hard constraint, which limits the solution space to the regions where the constraint is guaranteed. Thus, the predictive performance of DeepSaDe models can be worse than existing approaches which do not guarantee constraint satisfaction. For UC1, the predictive performance

UC	Approach	Constraint	AdI($\delta = 0.1$)	Accuracy/MSE	Training Time (sec)
UC1	DeepSaDe REG	100\pm0	0\pm0	*38.36 \pm 4.59	102341 \pm 40198
		93.50 \pm 2.01	0.97 \pm 0.003	*30.50\pm6.49	227 \pm 70
UC2	DeepSaDe SL SBR	100\pm0	0\pm0	80.04 \pm 4.29	447 \pm 105
		100\pm0	0.002 \pm 0.006	80.17\pm3.88	45 \pm 30
		100\pm0	0.002 \pm 0.004	80.04 \pm 3.95	45 \pm 27
UC3	DeepSaDe SL SBR	100\pm0	0\pm0	80.11 \pm 4.99	6580 \pm 1915
		99.97 \pm 0.03	0.42 \pm 0.28	82.53\pm3.58	101 \pm 45
		99.97 \pm 0.05	0.29 \pm 0.10	76.51 \pm 7.50	196 \pm 68

Table 1: Results: UC1, UC2 & UC3 (*MSE, lower MSE value is better)

UC	Approach	Constraint	Coherent	Flattened	Jaccard	Training Time (sec)
UC4	DeepSaDe FFN	100\pm0	6.62 \pm 1.72	78.52 \pm 1.73	62.97 \pm 2.28	227928 \pm 30559
		88.00 \pm 3.26	23.94\pm4.25	85.81\pm1.18	71.55\pm2.40	3215 \pm 2663
UC5	DeepSaDe FFN SL	100\pm0	11.08\pm2.61	67.17 \pm 1.48	25.94\pm2.89	17586 \pm 5074
		0.04 \pm 0.15	0.01 \pm 0.04	75.69\pm0.15	13.04 \pm 1.06	48 \pm 10
		100\pm0	4.06 \pm 3.33	63.16 \pm 2.62	18.08 \pm 3.33	298 \pm 110

Table 2: Results: UC4 & UC5

of DeepSaDe is slightly worse than REG, while the difference is not statistically significant for UC2 & UC3. The performance of DeepSaDe is worse for UC4 on all prediction metrics compared to FFN. For UC5, SL regularization with a high value of λ allows for satisfying all the constraints on the test set but performs much worse than DeepSaDe. To study this further, we plot the prediction loss (cross-entropy) for SL models on the test set for various λ between 0 and 0.9, averaged over 5 folds, in Figure 2. For comparison, the average loss for DeepSaDe is also plotted. DeepSaDe achieves constraint satisfaction in addition to having better performance than SL with a high λ . High regularization makes the prediction loss insignificant compared to the regularization loss, leading to worse predictive performance.




Figure 2: Test loss for UC5

In DeepSaDe, at every iteration, the solver tries to satisfy as many soft constraints as possible in addition to satisfying the domain constraints. DeepSaDe, thus, is a more stable learner compared to regularization with high λ . This answers **Q2**. Although the constraint satisfaction with DeepSaDe comes at the cost of predictive performance in some cases, in applications where constraints are crucial, like in

safety-critical domains, this may be acceptable.

Training Time: DeepSaDe requires between 10 to 500 times more time than baselines across use cases. This is because DeepSaDe solves a MaxSMT problem at each iteration, which makes the training slower than the numerical updates in the baselines. This positively answers **Q3**. In applications where constraints are imperative, training time is less relevant, e.g., a network trained over a week and guarantees safety is still more valuable to the autonomous vehicle compared to one trained for a few hours but cannot do so. Our work is a starting point of such an approach that combines the exact solving of quantifiers with gradient-based learning. With further research into solver technology, it can be made more scalable. Additionally, possible modifications to improve the efficiency of DeepSaDe include using an incomplete MaxSMT approach like stochastic local search (Morgado et al. 2013) instead of a complete one like Fu-Malik, and using compact latent bounds (Wang et al. 2018).

7 Conclusion

We introduced DeepSaDe, a method for training feed-forward neural networks to satisfy a range of constraints through a combination of a satisfiability framework and gradient-based optimization. DeepSaDe effectively addresses various ML tasks and offers adaptable constraint representation, although sometimes with a performance trade-off. It enforces any constraint in an SMT formula, feasible for the solver. While Z3 is our current solver, our framework is adaptable to any MaxSMT solver handling universally quantified constraints. We believe that evolving solver capabilities will allow DeepSaDe to handle more complex constraints. Extension of DeepSaDe to other architectures (e.g. CNNs) is left for future work.

Acknowledgments

This research was jointly funded by the Flemish Government (AI Research Program) and the Research Foundation - Flanders under EOS No. 30992574.

References

- Berg, O. J.; Hyttinen, A. J.; Järvisalo, M. J.; et al. 2019. Applications of MaxSAT in data analysis. *Pragmatics of SAT*.
- Bertsimas, D.; and Dunn, J. 2017. Optimal classification trees. *Machine Learning*, 106(7): 1039–1082.
- Bunel, R. R.; Turkaslan, I.; Torr, P.; Kohli, P.; and Mudigonda, P. K. 2018. A unified view of piecewise linear neural network verification. *Advances in Neural Information Processing Systems*, 31.
- Chavira, M.; and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *Artificial Intelligence*, 172(6-7): 772–799.
- Cussens, J. 2012. Bayesian network learning by compiling to weighted MAX-SAT. *arXiv preprint arXiv:1206.3244*.
- Demirović, E.; Lukina, A.; Hebrard, E.; Chan, J.; Bailey, J.; Leckie, C.; Ramamohanarao, K.; and Stuckey, P. J. 2022. MurTree: Optimal Decision Trees via Dynamic Programming and Search. *Journal of Machine Learning Research*, 23(26): 1–47.
- Diligenti, M.; Gori, M.; and Sacca, C. 2017. Semantic-based regularization for learning and inference. *Artificial Intelligence*, 244: 143–165.
- Fischer, M.; Balunovic, M.; Drachler-Cohen, D.; Gehr, T.; Zhang, C.; and Vechev, M. 2019. DI2: Training and querying neural networks with logic. In *International Conference on Machine Learning*, 1931–1941. PMLR.
- Fu, Z.; and Malik, S. 2006. On solving the partial MAX-SAT problem. In *International Conference on Theory and Applications of Satisfiability Testing*, 252–265. Springer.
- Goodfellow, I.; Bengio, Y.; Courville, A.; and Bengio, Y. 2016. *Deep learning*, volume 1. MIT Press.
- Gowal, S.; Dvijotham, K.; Stanforth, R.; Bunel, R.; Qin, C.; Uesato, J.; Arandjelovic, R.; Mann, T.; and Kohli, P. 2018. On the effectiveness of interval bound propagation for training verifiably robust models. *arXiv preprint arXiv:1810.12715*.
- Goyal, K.; Dumancic, S.; and Blockeel, H. 2023. Sade: Learning models that provably satisfy domain constraints. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2022, Grenoble, France, September 19–23, 2022, Proceedings, Part V*, 410–425. Springer.
- Günlük, O.; Kalagnanam, J.; Li, M.; Menickelly, M.; and Scheinberg, K. 2021. Optimal decision trees for categorical data via integer programming. *Journal of global optimization*, 81(1): 233–260.
- Gupta, M.; Cotter, A.; Pfeifer, J.; Voevodski, K.; Canini, K.; Mangylov, A.; Moczydlowski, W.; and Van Esbroeck, A. 2016. Monotonic calibrated interpolated look-up tables. *The Journal of Machine Learning Research*, 17(1): 3790–3836.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, 1026–1034.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hoernle, N.; Karampatsis, R. M.; Belle, V.; and Gal, K. 2022. Multiplexnet: Towards fully satisfied logical constraints in neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 5700–5709.
- Hu, Z.; Ma, X.; Liu, Z.; Hovy, E.; and Xing, E. 2016. Harnessing deep neural networks with logic rules. *arXiv preprint*.
- Ignatiev, A.; Lam, E.; Stuckey, P. J.; and Marques-Silva, J. 2021. A scalable two stage approach to computing optimal decision sets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3806–3814.
- Katz, G.; Barrett, C.; Dill, D. L.; Julian, K.; and Kochenderfer, M. J. 2017. Reluplex: An efficient SMT solver for verifying deep neural networks. In *International conference on computer aided verification*, 97–117. Springer.
- Katz, G.; Huang, D. A.; Ibeling, D.; Julian, K.; Lazarus, C.; Lim, R.; Shah, P.; Thakoor, S.; Wu, H.; Zeljić, A.; et al. 2019. The marabou framework for verification and analysis of deep neural networks. In *International Conference on Computer Aided Verification*, 443–452. Springer.
- Leino, K.; Fromherz, A.; Mangal, R.; Fredrikson, M.; Parno, B.; and Păsăreanu, C. 2022. Self-correcting Neural Networks for Safe Classification. In *International Workshop on Numerical Software Verification, Workshop on Formal Methods for ML-Enabled Autonomous Systems*, 96–130. Springer.
- Lin, X.; Zhu, H.; Samanta, R.; and Jagannathan, S. 2020. ART: abstraction refinement-guided training for provably correct neural networks. In *# PLACEHOLDER_PARENT_METADATA_VALUE#*, volume 1, 148–157. TU Wien Academic Press.
- Malioutov, D.; and Meel, K. S. 2018. Mlic: A maxsat-based framework for learning interpretable classification rules. In *CP*.
- Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. *Advances in Neural Information Processing Systems*, 31.
- Morgado, A.; Heras, F.; Liffiton, M.; Planes, J.; and Marques-Silva, J. 2013. Iterative and core-guided MaxSAT solving: A survey and assessment. *Constraints*, 18(4): 478–534.
- Moura, L. d.; and Bjørner, N. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, 337–340. Springer.
- Nethercote, N.; Stuckey, P. J.; Becket, R.; Brand, S.; Duck, G. J.; and Tack, G. 2007. MiniZinc: Towards a standard

- CP modelling language. In *International Conference on Principles and Practice of Constraint Programming*, 534. Springer.
- Serafini, L.; and Garcez, A. d. 2016. Logic tensor networks: Deep learning and logical reasoning from data and knowledge. *arXiv preprint arXiv:1606.04422*.
- Sivaraman, A.; Farnadi, G.; Millstein, T.; and Van den Broeck, G. 2020. Counterexample-guided learning of monotonic neural networks. *Advances in Neural Information Processing Systems*, 33: 11936–11948.
- Stewart, R.; and Ermon, S. 2017. Label-free supervision of neural networks with physics and domain knowledge. In *AAAI*.
- Sunaga, T. 1958. Theory of interval algebra and its applications to numerical analysis. *RAAG Memoirs, 2-Misc. II*, 547–564.
- Verwer, S.; and Zhang, Y. 2019. Learning optimal classification trees using a binary linear program formulation. In *Proceedings of the AAAI conference on artificial intelligence*, 1625–1632.
- Wang, S.; Pei, K.; Whitehouse, J.; Yang, J.; and Jana, S. 2018. Formal security analysis of neural networks using symbolic intervals. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, 1599–1614.
- Xu, J.; Zhang, Z.; Friedman, T.; Liang, Y.; and Broeck, G. 2018. A semantic loss function for deep learning with symbolic knowledge. In *International conference on machine learning*, 5502–5511. PMLR.
- You, S.; Ding, D.; Canini, K.; Pfeifer, J.; and Gupta, M. 2017. Deep lattice networks and partial monotonic functions. *Advances in neural information processing systems*, 30.
- Yu, J.; Ignatiev, A.; Stuckey, P. J.; and Le Bodic, P. 2021. Learning Optimal Decision Sets and Lists with SAT. *Journal of Artificial Intelligence Research*, 72: 1251–1279.
- Zhang, H.; Chen, H.; Xiao, C.; Gowal, S.; Stanforth, R.; Li, B.; Boning, D.; and Hsieh, C.-J. 2019. Towards stable and efficient training of verifiably robust neural networks. *arXiv preprint arXiv:1906.06316*.