Faster Verified Explanations for Neural Networks

Alessandro De Palma[1]
London School of Economics and Political Science, UK

Greta Dolcetti
Ca' Foscari University of Venice, Italy

Caterina Urban
Inria & ENS | PSL, Paris, France

Abstract

Verified explanations are a principled way to explain the decisions taken by neural networks, which are otherwise black-box in nature. However, these techniques face significant scalability challenges, as they require multiple calls to neural network verifiers, each of them with an exponential worst-case complexity. We present FaVeX, a novel algorithm to compute verified explanations. FaVeX accelerates the computation by dynamically combining batch and sequential processing of input features, and by reusing information from previous queries, both when proving invariances with respect to certain input features, and when searching for feature assignments altering the prediction. Furthermore, we present a novel and hierarchical definition of verified explanations, termed verifier-optimal robust explanations, that explicitly factors the incompleteness of network verifiers within the explanation. Our comprehensive experimental evaluation demonstrates the superior scalability of both FaVeX, and of verifier-optimal robust explanations, which together can produce meaningful formal explanation on networks with hundreds of thousands of non-linear activations.

2012 ACM Subject Classification Theory of computation → Program analysis; Theory of computation → Abstraction; Mathematics of computing → Continuous optimization

Keywords and phrases Verified Explanations, eXplainable Artificial Intelligence (XAI), Local Robustness, Neural Network Verification, Static Analysis

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2026.3

Related Version *Full Version*: `https://arxiv.org/abs/2512.00164`

Supplementary Material
Software (Source Code): `https://github.com/alessandrodepalma/favex` [13]

Funding This work was partially supported by the SAIF project, funded by the "France 2030" government investment plan managed by the French National Research Agency, under the reference ANR-23-PEIA-0006, as well as the ForML project, funded by the French National Research Agency, under the reference ANR-23-CE25-0009.

Acknowledgements The authors are grateful to the CLEPS infrastructure from Inria Paris for providing resources and support.

1 Introduction

Machine learning models have demonstrated remarkable performance across a wide range of tasks, from image classification to natural language processing. Yet, as these models are increasingly deployed in safety-critical domains – such as finance, transportation, and even healthcare – concerns about trust, accountability, and robustness have become paramount. In such contexts, explainability is essential: users and auditors must understand *why* a model produces a particular prediction, and whether that prediction remains stable under small, semantically meaningful perturbations of the input [20, 15].

[1] Work partly carried out at Inria & ENS & PSL, France.

(a) Robust explanation. **(b)** Verifier-optimal robust explanation.

Figure 1 Standard (a) and verifier-optimal (b) robust explanations ($\epsilon = 0.25$) of the fifth image of the MNIST test set for a state-of-the-art network from the certified training literature [12]. *Counterfactuals* are highlighted in *yellow*, *unknowns* in *blue*, and *invariants* are not highlighted. While the robust explanation is too large to be meaningful, the counterfactuals in (b) point to an interpretable digit transformation.

A promising line of research in eXplainable Artificial Intelligence addresses this need through *verified explanations*, e.g., [43, 44], among others, which aim to formally certify which input features are responsible for a model prediction by leveraging the theory and tools of machine learning verification [1, 61]. Thus, verified explanations aim to provide rigorous guarantees, in contrast to heuristic or gradient-based methods, e.g., [23, 40, 50]. Among verified explanations, *(optimal) robust explanations* [28, 35, 66, 26] have emerged as principled and mathematically grounded notions linking local robustness to the minimality of feature sets preserving a model prediction: they identify subsets of input features that, if left unchanged, cannot alter the predicted outcome. Formally, this corresponds to establishing local robustness of the model with respect to perturbations on the remaining input features.

However, despite their theoretical appeal, optimal robust explanations are difficult to scale beyond small neural networks with tens of non-linear activations or low-dimensional input spaces [66, 65]. Indeed, the computation of optimal robust explanations requires repeatedly querying neural network verifiers, which solve NP-hard problems [33] and often exhibit their exponential worst-case complexity. In practice, even finding a single counterfactual – a necessary step to ensure the *optimality* of an explanation – can quickly become infeasible for state-of-the-art neural networks with hundreds of thousands of non-linear activations (we demonstrate and discuss this in detail in Section 7.2).

Crucially, the definition of optimal robust explanations assumes that verification is complete, i.e., always either proving robustness or producing a counterfactual. This assumption is unrealistic in practice because, due to their exponential worst-case complexity, modern verifiers have to routinely enforce timeouts to handle large models efficiently, de facto making compromises about completeness to gain better performance [5, 33, 59]. As a consequence, although optimal explanations are a useful semantic ideal, they are rarely practically achievable for modern architectures. For instance, Figure 1a shows the robust (but not optimal) explanation that can be computed in practice for the classification of the fifth image in the MNIST testing data set by a state-of-the-art convolutional neural network (with roughly $230k$ ReLUs) from the certified training literature [12]. Notably, the explanation (i.e., the pixels that are highlighted in blue) covers a large portion of the image, making it too coarse to provide meaningful insights in practical settings. On the other hand, the counterfactuals (highlighted in yellow) directly point to a highly-interpretable digit transformation.

Contributions. We address the aforementioned challenges through three main contributions:

1. **Verifier-Optimal Robust Explanations (Section 4).** We introduce a new and practically achievable counterpart of optimal robust explanations – *verifier-optimal robust explanations* – that explicitly incorporates the behavior of a given verifier and its de

facto incompleteness. Our notion partitions the input features into: (i) *invariants*, for which the verifier proves robustness, (ii) *counterfactuals*, for which the verifier identifies perturbations causing misclassification, and (iii) *unknowns*, where verification remains inconclusive. This viewpoint better reflects what can be guaranteed in practice and enables the discovery of counterfactuals at a scale unattainable with prior definitions. Figure 1b shows the verifier-optimal robust explanation computed by our approach for the same MNIST image and convolutional model of Figure 1a. In this case, the explanation is *hierarchical*, with the most informative pixels (i.e., the counterfactuals, highlighted in yellow) covering only a small portion of the image. Importantly, the definition of verifier-optimal robust explanation is a generalization of optimal robust explanations: when the verifier is complete, the set of unknowns is empty and verifier-optimal robust explanations coincide exactly with optimal robust explanations (Lemma 4).

2. **FaVeX: A Fast Algorithm for Computing Verified Explanations (Section 5)**. We propose a new algorithm that substantially accelerates the computation of both standard and verifier-optimal explanations. FaVeX incorporates three key innovations:
   - **(a)** a hybrid query processing strategy that dynamically combines batch-based processing with binary search and sequential feature evaluation (Section 5.1);
   - **(b)** branch reuse within branch-and-bound verification to leverage information from previous verification queries (Section 5.2); and
   - **(c)** a restricted-space adversarial attack to accelerate counterfactual search (Section 5.3).

   We also introduce a traversal strategy aligned with the verifier's logit-difference objective (Algorithm 4).
3. **Implementation and Experimental Evaluation (Sections 6 and 7)**. We implemented FaVeX using the PyTorch deep learning library, employing the OVAL branch-and-bound framework [6, 4, 11, 10] and its α-β-CROWN implementation [67, 63] as the backbone for our verifier. We evaluate it on both small networks similar to those employed in the verified explanation literature [66, 65], and on significantly larger state-of-the-art certifiably-robust neural networks. The results demonstrate that:
   - **(i)** FaVeX significantly reduces the time required to compute standard robust explanations on small fully-connected networks;
   - **(ii)** verifier-optimal robust explanations can be computed efficiently even on larger convolutional networks; and
   - **(iii)** FaVeX finds counterfactuals even for networks with hundreds of thousands of activations, making formal explanations practical and meaningful at scale.

2 Standard vs. Verifier-Optimal Robust Explanations

In this section, through a simple example[2], we provide an informal overview of the difference between our verifier-optimal robust explanations and standard robust explanations.

Let us consider the classifier with 9-dimensional input in Figure 2 and let $\mathbf{x} = (x_1, \ldots, x_9) = (1.0, 0.7, 0.7, 0.2, 0.8, 0.4, 0.7, 0.3, 0.2)$ be the input vector for which we want to compute an explanation. For simplicity, we employ a basic deletion-based algorithm [66, 26] that processes input features indexes in sequential order from 1 to 9. Each step introduces ϵ-bounded ℓ_∞ perturbations (i.e., such that the maximum absolute value difference for any

[2] A self-contained script to reproduce this example is available in our source code repository at: `https://github.com/alessandrodepalma/favex/blob/main/ecoop2026.py`

[3] `https://archive.ics.uci.edu/dataset/15/breast+cancer+wisconsin+original`

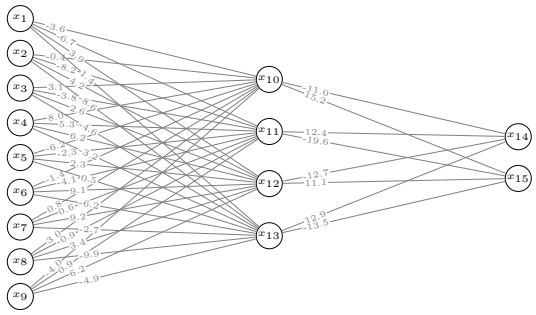

Figure 2 Fully connected feed-forward neural network classifier with ReLU activations, trained on the original Wisconsin Breast Cancer Database[3].

feature is at most ϵ) to the input feature under analysis and selected previous features (depending on the explanation definition), while all remaining input features remain unchanged. Let us fix $\epsilon = 0.6$ for this example. A verifier is then queried to determine whether this perturbation is robust (i.e., the model's prediction remains the same), a counterexample is found (i.e., applying the perturbation causes the model to change its prediction), or a timeout is reached (i.e., the query cannot be resolved within the time limit).

Standard Robust Explanations. Figure 3 shows how the analysis proceeds to compute a standard robust (but, in practice, not optimal) explanation using branch-and-bound verification with CROWN / DeepPoly [69, 57] (cf. Section 3.1). First, the verifier is queried with the ℓ_∞ perturbation applied only to x_1. In this case, robustness is verified: the feature index is added to the invariants set $\mathcal{R}_\mathbf{x}$ (initially empty). Next, the perturbation is applied to the features indexed by $\mathcal{R}_\mathbf{x}$ (currently just x_1) and to the new feature under analysis, x_2. Robustness is again verified: this feature index is added to $\mathcal{R}_\mathbf{x}$. The analysis then proceeds with x_3, x_4, x_5, finding all of them to be invariants. Next, for x_6, the verifier finds a counterexample. This feature is therefore considered part of the explanation and will not be perturbed in future steps. The analysis continues with x_7. When x_7 is perturbed (along with the invariants x_1, x_2, x_3, x_4, and x_5), the verifier times out. Since the definition of optimal robust explanation does not account for the incompleteness caused by timeouts, the feature is considered part of the explanation from this point on. The computation continues until all input features have been processed: another timeout is encountered for x_9, while x_8 is invariant. The explanation indiscriminately comprises features x_6, x_7, and x_9. However, given the verifier timeouts, x_7 and x_9 may actually be invariants.

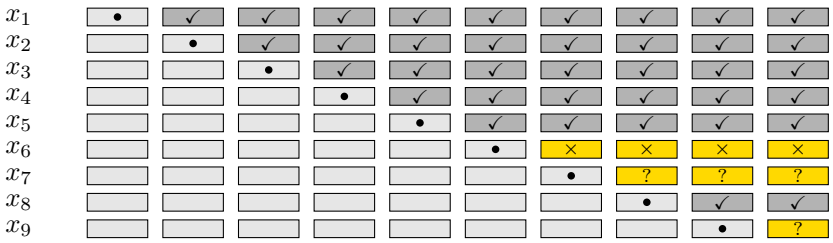

Figure 3 Example of computing a standard robust explanation on 9 features. The verification steps proceed from left to right; the dotted box indicates the feature under analysis, dark grey shows the invariants (✓). Yellow denotes the explanations, including both features for which a counterexample has been found (×), and those for which a timeout has been encountered (?). The final results of the analysis can be seen in the rightmost column.

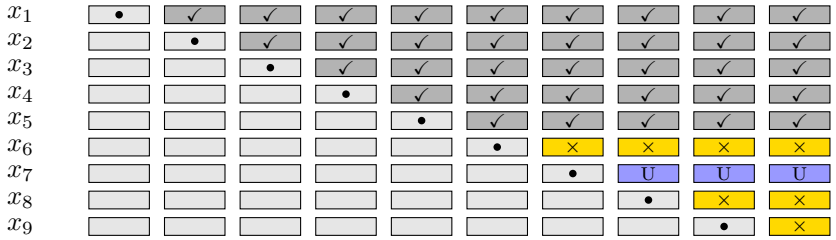

Figure 4 Example of computing a verifier-optimal robust explanation on 9 features. The verification steps proceed from left to right; the dotted box indicates the feature under analysis, dark grey shows the invariants (✓), blue shows the unknowns (U), yellow shows the explanations for which a counterexample has been found (×). The final results of the analysis can be seen in the rightmost column.

Verifier-Optimal Robust Explanations. Figure 4 instead illustrates the computation of a verifier-optimal robust explanation. In this case the analysis explicitly accounts for verifier limitations due to timeouts. Thus, rather than considering x_7 part of the explanation, the analysis adds the feature index to the unknowns set $\mathcal{U}_{\mathbf{x}}$ (initially empty). The analysis then proceeds by allowing perturbations also to the features indexed by $\mathcal{U}_{\mathbf{x}}$. This way, a counterexample is found for both x_9 and x_8 (which was instead invariant in Figure 3). This can occur because the considered perturbation region is effectively larger, as it also includes the features indexed by $\mathcal{U}_{\mathbf{x}}$. (This effect can also be observed in practice beyond this illustrative example: for instance, Figure 1b has 4 more invariant pixels than Figure 1a.) The explanation, containing the features x_6, x_7, x_8, x_9, is larger than the one in Figure 3 but it is *hierarchical*: features x_6, x_8, x_9 (counterfactuals) are more important to the classification outcome, while x_7 (unknown) is less important.

3 Background

In the following, we will denote vectors by boldface lowercase letters ($\mathbf{x} \in \mathbb{R}^d$), and use subscripts to index their entries (for instance, $\mathbf{x}_i$ denotes the i-th entry). Furthermore, we will denote sets by calligraphic uppercase letters (e.g, $\mathcal{A}$), and use $\mathbf{x}_{\mathcal{A}}$ to denote the vector containing the entries of $\mathbf{x}$ indexed by the elements of $\mathcal{A}$.

3.1 Neural Network Verification

Neural Networks. A *neural network* classifier is a function $f : \mathbb{R}^d \to \mathbb{R}^k$ mapping a d-dimensional input vector of *features* $\mathbf{x}$ to a k-dimensional output vector $f(\mathbf{x})$ of *logits*, which determine the network's classification output $y_f : \mathbb{R}^k \to \{1, \dots, k\}$, defined as $y_f(\mathbf{x}) \stackrel{\text{def}}{=} \operatorname{argmax}_i f(\mathbf{x})_i$. We here focus on neural networks with rectified linear unit (ReLU) activations, which consist of a sequential composition $f = f_l \circ \cdots \circ f_1$ of l layers, where each layer $f_i \colon \mathbb{R}^m \to \mathbb{R}^n$, is either an affine transformation $f_i(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, with $\mathbf{A} \in \mathbb{R}^{n \times m}$, $\mathbf{b} \in \mathbb{R}^n$, or an entry-wise application of the ReLU activation function, $\text{ReLU}(x) \stackrel{\text{def}}{=} \max(0, x)$; the final layer f_l is an affine transformation.

Local Robustness. Neural network verification aims to provide formal guarantees about neural network behavior. In particular, verifying neural network *safety*, involves proving that all network outputs corresponding to inputs satisfying a given input property also satisfy

a specified output property. A widely studied safety property is *local robustness*, which ensures that small (up to a chosen threshold) perturbations of an input do not change the classification output of a neural network. Formally, a neural network classifier $f : \mathbb{R}^d \to \mathbb{R}^k$ is said to be locally robust on input $\mathbf{x} \in \mathbb{R}^d$ if, given a set of allowed perturbations $C(\mathbf{x}) \subseteq \mathbb{R}^d$, for all $\mathbf{x}' \in \mathbb{R}^d$ we have:

$$\mathbf{x}' \in C(\mathbf{x}) \Rightarrow y_f(\mathbf{x}') = y_f(\mathbf{x})$$

that is, the classification output of the neural network is invariant to perturbations in $C(\mathbf{x})$. In the following, we focus on perturbations within the ℓ_∞-ball with radius $\epsilon > 0$ centered at $\mathbf{x}$, i.e., $C(\mathbf{x}) = B^\epsilon(\mathbf{x}) \stackrel{\text{def}}{=} \{\mathbf{x}' \in \mathbb{R}^d \mid \|\mathbf{x}' - \mathbf{x}\|_\infty \leq \epsilon\}$, where $\|\mathbf{v}\|_\infty \stackrel{\text{def}}{=} \max_i |v_i|$ denotes the ℓ_∞ norm of a vector $\mathbf{v}$. In some cases, we are interested in perturbations that affect only a subset of the input features. Let $\mathcal{A} \subseteq \{1, \ldots, d\}$ denote the indices of input features that may be perturbed. We will write $\overline{\mathcal{A}} := \{1, \ldots, d\} \setminus \mathcal{A}$ for the complement of $\mathcal{A}$ under $\{1, \ldots, d\}$. We define the set of allowed perturbations restricted to $\mathcal{A}$ as $B^\epsilon_{\mathcal{A}}(\mathbf{x}) \stackrel{\text{def}}{=} \{\mathbf{x}' \in \mathbb{R}^d \mid \mathbf{x}'_{\overline{\mathcal{A}}} = \mathbf{x}_{\overline{\mathcal{A}}} \wedge \|\mathbf{x}'_{\mathcal{A}} - \mathbf{x}_{\mathcal{A}}\|_\infty \leq \epsilon\}$.

We define a *robustness query* as a tuple $(\mathbf{x}, \mathcal{A}, \epsilon, f)$, where $\mathbf{x}$ denotes the input vector, $\mathcal{A}$ the set of the indices of the perturbed input features, ϵ the perturbation radius, and f the employed network. Specifically, $(\mathbf{x}, \mathcal{A}, \epsilon, f)$ amounts to assessing local robustness of f on $\mathbf{x}$ given the set of allowed perturbations $B^\epsilon_{\mathcal{A}}(\mathbf{x})$. Let $\mathcal{Q}$ be the set of all robustness queries. A *neural network verifier* $v : \mathcal{Q} \to \{-1, 0, 1\}$ is a function taking a robustness query as input, and returning: 1 if local robustness is proved; -1 if a *counterfactual* is found, that is, an input $\mathbf{x}' \in B^\epsilon_{\mathcal{A}}(\mathbf{x})$ such that $y_f(\mathbf{x}') \neq y_f(\mathbf{x})$; and 0 in case the verification is inconclusive. A *complete* verifier $\overline{v} : \mathcal{Q} \to \{-1, 1\}$ never returns 0, it always either verifies the query if it holds (returns 1), or finds a counterfactual (returns -1).

Internally, a neural network verifier v leverages an *analyzer* $a : \mathcal{Q} \to \mathbb{R}$ which, given a robustness query $(\mathbf{x}, \mathcal{A}, \epsilon, f)$, computes lower and upper bound values of each logit in $f(\mathbf{x})$ for any input vector $\mathbf{x}' \in B^\epsilon_{\mathcal{A}}(\mathbf{x})$. To do so it can leverage different abstractions such as IBP [18, 19], CROWN / DeepPoly [69, 57], or α-CROWN [67], among others [38, 47, 56, 55, 62, etc.]. Based on these bounds, a returns the lower bound value of the *worst-case logit difference*, the minimum difference between the logit corresponding to the true class $f(\mathbf{x})_{y_f(\mathbf{x})}$ and the logit corresponding to any other class: $\min_{i \neq y_f(\mathbf{x})}(f(\mathbf{x})_{y_f(\mathbf{x})} - f(\mathbf{x})_i)$. This lower bound is central to verification: if strictly positive, the logit corresponding to the true class remains larger than all others for every perturbed input vector in $B^\epsilon_{\mathcal{A}}(\mathbf{x})$, thus local robustness of f on x is verified (the verifier v can return 1).

Branch-and-Bound. Complete neural network verifiers can be interpreted under the lens of the *branch-and-bound* paradigm [6], explicitly adopted by state-of-the-art complete neural network verifiers (e.g., [63, 70]), which recursively partitions the verification problem into more tractable subproblems.

Let $\hat{\mathcal{Q}} \stackrel{\text{def}}{=} \{(\mathbf{x}, \mathcal{A}, \epsilon, f, C) \mid (\mathbf{x}, \mathcal{A}, \epsilon, f) \in \mathcal{Q}, C \in \mathcal{P}(\mathcal{C})\}$ be an extension of $\mathcal{Q}$ where each robustness query $(\mathbf{x}, \mathcal{A}, \epsilon, f) \in \mathcal{Q}$ is augmented by a set of splitting constraints $C \in \mathcal{P}(\mathcal{C})$, where $\mathcal{C}$ denotes the set of all possible splitting constraints and $\mathcal{P}(\mathcal{C})$ denotes its powerset. We refer to elements of $\hat{\mathcal{Q}}$ as *robustness subproblems*. We assume that the analyzer $a \colon Q \to \mathbb{R}$ naturally accommodates splitting constraints. Therefore, from now on, we assume that analyzers directly operate on subproblems: $a \colon \hat{Q} \to \mathbb{R}$. For a subproblem $Q = (\mathbf{x}, \mathcal{A}, \epsilon, f, C) \in \hat{\mathcal{Q}}$, we write Q_C to denote its associated constraint set C. Branch-and-bound partitions a given robustness subproblem $Q \in \hat{\mathcal{Q}}$ by growing its set of constraints Q_C. This is typically done by either partitioning the set of allowed perturbations (input splitting [2, 62]), or activations: in the following, we focus on the latter, termed *ReLU splitting* [5, 11, 25] for the

Algorithm 1 Branch-and-Bound.

```
1: function BaB(a, (x, A, ϵ, f))                    ▷ a: Q̂ → ℝ, (x, A, ϵ, f) ∈ Q
2:     unresolved ← {(x, A, ϵ, f, ∅)}                ▷ (x, A, ϵ, f, ∅) ∈ Q̂
3:     for Q ∈ unresolved do
4:         if CEX(x, A, ϵ, f) then return −1
5:         unresolved ← unresolved \ {Q}
6:         result ← a(Q)                             ▷ Bounding step
7:         if result ≤ 0 then                        ▷ Branching step
8:             Q1, Q2 ← SPLIT(Q)
9:             unresolved ← unresolved ∪ {Q1, Q2}
10:    return 1
```

networks under consideration, owing to its higher efficacy on high-dimensional input feature spaces. That is, we assume that constraints in $\mathcal{C}$ determine the sign of the pre-activation value of a ReLU: let $x_{i,j} = \max(0, \hat{x}_{i,j})$ denote a ReLU at the i-th layer and j-index of a neural network classifier, where $\hat{x}_{i,j}$ represents the pre-activation value input to the ReLU; the constraint $\hat{x}_{i,j} \geq 0$ determines that the ReLU behaves as the identify function, while the constraint $\hat{x}_{i,j} < 0$ determines that the ReLU behaves as the constant function equal to zero.

Branch-and-bound verification is illustrated by Algorithm 1. The procedure maintains a list of unresolved robustness subproblems, initially containing the given robustness query augmented with an empty set of constraints (cf. Line 2). Each unresolved subproblem Q spawns a counterexample search (cf. Line 4). If the search is successful, a counterfactual is found and the procedure terminates immediately. Otherwise, the subproblem Q is given to the analyzer a in the *bounding step* (cf. Line 6). If $v(Q) \leq 0$, the verification is inconclusive and the algorithm performs the *branching step* (cf. Line 7): the unresolved subproblem Q is split into further subproblems Q_1 and Q_2 (cf. Line 8) which are added back to the list of unresolved subproblems (cf. Line 9). The splitting is done according to the partitioning strategy; in case of ReLU splitting, the branching refines the subproblem $Q = (\mathbf{x}, \mathcal{A}, \epsilon, f, C)$ by selecting a ReLU whose pre-activation value $\hat{x}_{i,j}$ is not yet constrained by C and partitioning the search space according to its possible sign status, yielding $Q_1 = (\mathbf{x}, \mathcal{A}, \epsilon, f, C \cup \{\hat{x}_{i,j} \geq 0\})$ and $Q_2 = (\mathbf{x}, \mathcal{A}, \epsilon, f, C \cup \{\hat{x}_{i,j} < 0\})$. The branch-and-bound process continues until all unresolved subproblems have been resolved, in which case robustness is verified (cf. Line 10).

Due to the NP-hardness of neural network verification [33], the number of subproblems generated by branch-and-bound can grow exponentially in the worst case. Therefore, practical implementations typically enforce a timeout, terminating the search and returning an inconclusive answer once the allocated verification time has elapsed (see Algorithm 2, where the differences with respect to Algorithm 1 are highlighted in boxes). An alternative strategy is to cap the number of branching steps or the total processed subproblems [11].

3.2 Verified Explanations

Optimal Robust Explanations. Recent work [28, 43, 44] has established a tight connection between local robustness and explainability of machine learning classifiers. In particular, *abductive explanations* [27, 46], *prime implicants* [53], and *sufficient reasons* [8, 9] characterize the minimal subset of input features that are responsible for a classifier prediction, in the sense that any perturbation on the rest of the input features will never change the classification output. Restricting perturbations to a set of allowed perturbations leads to *distance-restricted* or *robust explanations* [35, 66, 26]:

Algorithm 2 Branch-and-Bound with Timeout.

```
1:  function BaB(a, (x, A, ε, f), T)          ▷ a: Q̂ → ℝ, (x, A, ε, f) ∈ Q, T ∈ ℕ
2:      t1 ← TIME()
3:      unresolved ← {(x, A, ε, f, ∅)}        ▷ (x, A, ε, f, ∅) ∈ Q̂
4:      for Q ∈ unresolved do
5:          if CEX(x, A, ε, f) then return −1
6:          unresolved ← unresolved \ {Q}
7:          result ← a(Q)                     ▷ Bounding step
8:          if result ≤ 0 then                ▷ Branching step
9:              t2 ← TIME()
10:             if t2 − t1 < T then
11:                 Q1, Q2 ← SPLIT(Q)
12:                 unresolved ← unresolved ∪ {Q1, Q2}
13:             else return 0
14:     return 1
```

▶ **Definition 1** (Robust Explanation). *Given a classifier* $f\colon \mathbb{R}^d \to \mathbb{R}^k$*, an input vector* $\mathbf{x} \in \mathbb{R}^d$*, and a perturbation radius* $\epsilon > 0$*, a* robust explanation *is a subset of the indexes of the input features* $\mathcal{E}_\mathbf{x} \subseteq \{1, \ldots, d\}$ *such that* f *is locally robust on* $\mathbf{x}$ *to perturbations in* $B^\epsilon_{\overline{\mathcal{E}}_\mathbf{x}}(\mathbf{x})$*:*

$$\forall \mathbf{x}' \in B^\epsilon_{\overline{\mathcal{E}}_\mathbf{x}}(\mathbf{x})\colon y_f(\mathbf{x}') = y_f(\mathbf{x})$$

Owing to Definition 1, we will call $\overline{\mathcal{E}}_\mathbf{x}$ the *invariants* for $\mathbf{x}$. The right-most column of Figure 3 shows a robust explanation, which consists of the features for which a counterexample has been found (×) or the verifier timed out (?), while the other features are the invariants (✓). Note that, the set of all input features indexes $\{1, \ldots, d\}$ is a trivially robust explanation. More generally, many subsets of $\{1, \ldots, d\}$ can constitute a robust explanation. Of particular interest is an explanation containing no superfluous features [28, 66]:

▶ **Definition 2** (Optimal Explanation). *Given a classifier* $f\colon \mathbb{R}^d \to \mathbb{R}^k$*, an input vector* $\mathbf{x} \in \mathbb{R}^d$*, and a perturbation radius* $\epsilon > 0$*, a robust explanation* $\mathcal{E}_\mathbf{x}$ *is* optimal *if removing an index from* $\mathcal{E}_\mathbf{x}$ *would break local robustness:*

$$\forall\, i \in \mathcal{E}_\mathbf{x}\colon \exists\, \mathbf{x}^* \in B^\epsilon_{\overline{\mathcal{E}}_\mathbf{x} \cup \{i\}}(\mathbf{x})\colon y_f(\mathbf{x}^*) \neq y_f(\mathbf{x}) \tag{1}$$

Note that, based on Definition 2, establishing the optimality of a robust explanation $\mathcal{E}_\mathbf{x}$ entails obtaining a *counterfactual* $\mathbf{x}^*$ witnessing Equation (1) for each feature in $\mathcal{E}_\mathbf{x}$.

Computing Robust Explanations. Existing algorithms for computing (optimal) robust explanations exploit a local robustness verifier as an *oracle* for identifying invariants or finding counterfactuals. The simplest algorithms [66, 26] operate sequentially, mimicking the deletion-based algorithm [7] for computing minimal unsatisfiable subsets of logic formulas. Others [29, 65] implement dichotomic search [24] or adapt the quickXplain algorithm [32].

Despite these advances, scalability remains a challenge: computing optimal explanation is possible for neural network classifiers with 20-60 ReLUs [66, 65], while only robust but not necessarily optimal explanations can be computed for larger machine learning models with a large number of input features [66, 29]. In particular, in our experiments, we found that finding counterfactuals rapidly becomes infeasible. For instance, computing robust

explanations with a 60-second timeout per robustness query on state-of-the-art models (with ~230k ReLUs) from the certified training literature [12] finds no counterexamples despite multiple hours of computation per image (cf. Table 3 in Section 7).

To bridge this gap between the *ideal* notion of optimal robust explanations (Definition 2) and its *practical realizability*, we introduce in Section 4 the concept of *verifier-optimal robust explanations*, which explicitly takes into account the underlying verifier, thereby capturing what can be provably achieved in practice. Importantly, this definition enables finding counterfactuals even at larger scale. Building on this notion, Section 5 presents an efficient algorithm for computing (verifier-optimal) robust explanations.

4 Verifier-Optimal Robust Explanations

The definition of optimal robust explanation (Definition 2) implicitly requires that the underlying verifier is complete, i.e., always able to either verify invariance or find a counterfactual for a given robustness query. This makes the notion of optimal robust explanation inherently semantic: it characterizes what explanations *should* be in principle, regardless of the limitations of practical verifiers.

While this semantic notion provides the ideal target, in practice verifiers are incomplete and may fail to decide a robustness query. To account for such uncertainty, we introduce a *practically achievable* counterpart of optimal robust explanations, that captures what can be established using a given verifier v. Specifically, our verifier-aware definition partitions the input feature indexes into three disjoint sets:

1. the invariants $\mathcal{R}_{\mathbf{x}}$ (not highlighted in Figure 5): the feature indexes that are surely not part of the explanation, for which the verifier v proves robustness (returns 1);
2. the counterfactuals $\mathcal{C}_{\mathbf{x}}$ (highlighted in yellow in Figure 5): the feature indexes for which the verifier v finds a counterfactual (returns -1);
3. the unknowns $\mathcal{U}_{\mathbf{x}}$ (highlighted in blue in Figure 5): the feature indexes for which the verifier v remains inconclusive (returns 0);

Intuitively, $\mathcal{R}_{\mathbf{x}}$ corresponds to input features that do not affect the classification outcome, while the union $\mathcal{C}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}$ of the counterfactuals and the unknowns yields a *hierarchical explanation* based on the ability of the given verifier v to find counterfactuals.

We formally define our verifier-optimal explanations below.

▶ **Definition 3** (Verifier-Optimal Robust Explanation). *Given a classifier $f : \mathbb{R}^d \to \mathbb{R}^k$, an input vector $\mathbf{x} \in \mathbb{R}^d$, a perturbation radius $\epsilon > 0$, and a verifier $v : \mathcal{Q} \to \{-1, 0, 1\}$, a* v-optimal robust explanation *is the union $\mathcal{C}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}$ of disjoint subsets of the input feature indexes $\mathcal{C}_{\mathbf{x}} \subseteq \{1, \ldots, d\}$ and $\mathcal{U}_{\mathbf{x}} \subseteq \{1, \ldots, d\}$ such that:*

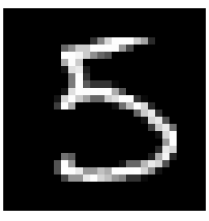

(a) Sixteenth image of the MNIST test set.

(b) Explanation for image (a) at $\epsilon = 0.25$.

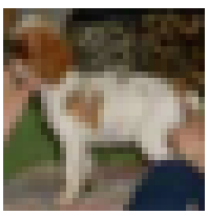

(c) Thirteenth image of the CIFAR-10 test set.

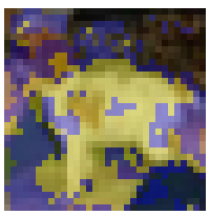

(d) Explanation for image (c) at $\epsilon = \frac{16}{255}$.

Figure 5 Examples of v-optimal robust explanation computed on state-of-the-art networks from the certified training literature [12].

1. $\mathcal{R}_{\mathbf{x}} \stackrel{def}{=} \overline{\mathcal{C}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}}$ *is the inclusion-maximal set of feature indexes such that the verifier* v *proves local robustness of* f *on* $\mathbf{x}$ *to perturbations in* $B^{\epsilon}_{\mathcal{R}_{\mathbf{x}}}(\mathbf{x})$*, i.e.,* $v((\mathbf{x}, \mathcal{R}_{\mathbf{x}}, \epsilon, f)) = 1$ *and, for all* $j \in \mathcal{C}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}$*,* $v((\mathbf{x}, \mathcal{R}_{\mathbf{x}} \cup \{j\}, \epsilon, f)) \neq 1$*;*
2. v *remains inconclusive when proving local robustness of* f *on* $\mathbf{x}$ *to perturbations in* $B^{\epsilon}_{\mathcal{R}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}}(\mathbf{x})$*, i.e.,* $v((\mathbf{x}, \mathcal{R}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}, \epsilon, f)) = 0$*;*
3. *the verifier* v *always finds counterfactuals when proving local robustness of* f *on* $\mathbf{x}$ *when perturbations are allowed on all feature indexed by* $\mathcal{R}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}}$ *and* any *feature indexed by* $\mathcal{C}_{\mathbf{x}}$*, i.e.,* $\forall\, i \in \mathcal{C}_{\mathbf{x}} \colon v((\mathbf{x}, \mathcal{R}_{\mathbf{x}} \cup \mathcal{U}_{\mathbf{x}} \cup \{i\}, \epsilon, f)) = -1$.

The right-most column of Figure 4 shows a verifier-optimal robust explanations, which consists of the counterfactuals in $\mathcal{C}_{\mathbf{x}}$ (×) and the unknowns in $\mathcal{U}_{\mathbf{x}}$ (U), while the other features are the invariants in $\mathcal{R}_{\mathbf{x}}$ (✓).

The first condition of Definition 3 implies that $\mathcal{U}_{\mathbf{x}} \cup \mathcal{C}_{\mathbf{x}}$ is a robust, yet in general not optimal, explanation. However, differently from Definitions 1 and 2, where perturbations are restricted to input features indexed by the invariants, in the second and third condition of Definition 3 perturbations are also allowed to all input features indexed by $\mathcal{U}_{\mathbf{x}}$. This increases the size of the perturbation space considered by the verifier, which in turn produces explanations that are *potentially larger* than those obtained when considering only perturbations on the invariants. On the other hand, we argue that these explanations are *more informative* since the input features indexed by $\mathcal{U}_{\mathbf{x}}$ do not produce counterfactuals individually, but counterfactuals found in combinations with these unknowns highlight clearly decisive input features for the classification outcome (the features indexed by $\mathcal{C}_{\mathbf{x}}$). In our experiments (cf. Tables 2 and 3 in Section 7), we found that also allowing perturbations on all input features indexed by $\mathcal{U}_{\mathbf{x}}$ is crucial to enable finding counterfactuals. Note that, Definition 2 is an instance of Definition 3 in which the verifier is complete $\overline{v} \colon \mathcal{Q} \to \{-1, 1\}$:

▶ **Lemma 4.** *Given a classifier* $f : \mathbb{R}^d \to \mathbb{R}^k$*, an input vector* $\mathbf{x} \in \mathbb{R}^d$*, a perturbation radius* $\epsilon > 0$*, and a* complete *verifier* $\overline{v} : \mathcal{Q} \to \{-1, 1\}$*, a* v-optimal robust explanation *(Definition 3) is an optimal robust explanation (Definition 2).*

Proof. Since the verifier is *complete*, it never returns 0. Thus, the second condition of Definition 3 never occurs. Hence, $\mathcal{U}_{\mathbf{x}} = \emptyset$ and Definition 3 coincides with Definition 2. ◀

5 FaVeX: Faster Verified Explanations

In this section, we introduce FaVeX, our efficient algorithm for computing (verifier-optimal) robust explanations. We first present the main algorithm (Section 5.1) and then describe two acceleration strategies: an approach to speed up verification by reusing prior branches in branch-and-bound verification (Section 5.2), and a technique to speed up the search for counterfactuals by means of adversarial attacks in a restricted space (Section 5.3).

5.1 Computing (Verifier-Optimal) Robust Explanations

FaVeX, shown in Algorithm 3, computes v-optimal robust explanations if the v-opt flag (cf. Line 1) is true, otherwise it computes robust explanations (cf. the parts highlighted in boxes in Algorithm 3). It maintains three sets $\mathcal{R}_{\mathbf{x}}$, $\mathcal{U}_{\mathbf{x}}$, and $\mathcal{C}_{\mathbf{x}}$, containing invariants, unknowns, and counterfactuals, all initially empty (cf. Line 2). It distinguishes between two kinds of robustness queries: *single queries*, in which perturbations are allowed on only one additional input feature besides the invariants in $\mathcal{R}_{\mathbf{x}}$ (cf. Lines 12 and 26) and, if the v-opt flag is true, the unknowns in $\mathcal{U}_{\mathbf{x}}$ (cf. Lines 11 and 25); and *batch queries*, in which

Algorithm 3 FaVeX.

```
1:  function FaVeX(v-opt, f, x, ε, a, A⃗, T)        ▷ v-opt ∈ {TRUE, FALSE}
                ▷ f: ℝ^d → ℝ^k, x ∈ ℝ^d, ε > 0, a: Q̂ → ℝ, A⃗ ∈ Sym({1,...,d}), T ∈ ℕ
2:      R_x, U_x, C_x ← ∅, ∅, ∅
3:      fallback ← FALSE
4:      batches ← {A⃗}
5:      for B ∈ batches do
6:          batches ← batches \ {B}
7:          if |B| > 1 then                          ▷ batch robustness query
8:              if fallback then                     ▷ fallback to sequential processing
9:                  for i ∈ B do
10:                     if v-opt then
11:                         result ← BaB(a, (x, R_x ∪ U_x ∪ {i}, ε, f), T)
12:                     else result ← BaB(a, (x, R_x ∪ {i}, ε, f), T)
13:                     if result == 1 then R_x ← R_x ∪ {i}
14:                     else if result = -1 then C_x ← C_x ∪ {i} else U_x ← U_x ∪ {i}
15:             else                                 ▷ binary search-based batch processing
16:                 if v-opt then
17:                     result ← BaB(a, (x, R_x ∪ U_x ∪ B, ε, f), T/10)
18:                 else result ← BaB(a, (x, R_x ∪ B, ε, f), T/10)
19:                 if result == 1 then R_x ← R_x ∪ B
20:                 else
21:                     A⃗_1, A⃗_2 ← halve(A⃗)
22:                     batches ← batches ∪ {A⃗_1, A⃗_2}
23:         else                                     ▷ single robustness query
24:             if v-opt then
25:                 result ← BaB(a, (x, R_x ∪ U_x ∪ B, ε, f), T)
26:             else result ← BaB(a, (x, R_x ∪ B, ε, f), T)
27:             if result == 1 then R_x ← R_x ∪ B
28:             else
29:                 fallback ← TRUE
30:                 if result = -1 then C_x ← C_x ∪ B else U_x ← U_x ∪ B
31:     return U_x, C_x
```

perturbations are allowed on multiple additional features simultaneously (cf. Line 17 and 18). For batch queries, the timeout T enforced on the underlying chosen verifier v (cf. Line 1) is reduced by a factor of 10 (cf. Line 17 and 18). The reduction factor can be a parameter of the algorithm but, in our experiments, we found this value to be effective in practice.

Similarly to [65], batch queries accelerate the identification of consecutive invariants by querying the verifier for multiple candidate input features at once, instead of making consecutive separate single queries for each of them. The search for batches of consecutive invariants is based on *binary search*. FaVeX takes as input a *traversal strategy* $\vec{\mathcal{A}}$, i.e., an ordered sequence of all the input feature indexes (cf. Line 1). The choice of traversal strategy matters for producing smaller explanations [66, 65]. We present our choice of traversal strategy later in this section. The first batch query allows perturbations on all input features indexed by $\vec{\mathcal{A}}$ (cf. Line 4). This query typically fails to verify (cf. Line 20), so $\vec{\mathcal{A}}$ is halved into two batches $\vec{\mathcal{A}}_1$ and $\vec{\mathcal{A}}_2$ (cf. Line 21) and binary search continues over them (cf. Line 22).

Algorithm 4 FaVeX Traversal Strategy.

```
1: function FaVeX-traversal(f, x, ϵ, a)        ▷ f: ℝ^d → ℝ^k, x ∈ ℝ^d, ϵ > 0, a: 𝒬 → ℝ
2:     S ← (0, ..., 0)                          ▷ S ∈ ℝ^d
3:     for i ∈ {1, ..., d} do
4:         S_i ← a((x, {i}, ϵ, f))
5:     return argsort(S, descending)
```

When the batch size is down to one (cf. Line 23), batch queries become de facto single queries, and they trigger a fallback to sequential feature processing (cf. Line 29) if the verifier disproves local robustness of f on $\mathbf{x}$ or times out (cf. Line 28). Once this occurs, unlike [65], we do not resume binary search over the next feature batches. In practice, we found that once batch queries have shrunk to size one and the verifier fails at this granularity, the likelihood of encountering further batches of consecutive invariants drops substantially. At this stage, continuing binary search adds overhead without providing benefits: it repeatedly queries increasingly small batches that are unlikely to yield new invariants, thereby worsening overall performance (cf. Tables 4-7 in Section 7). For this reason, after the fallback is triggered (cf. Line 8), FaVeX switches entirely to sequential feature processing (cf. Lines 9-14).

The sets $\mathcal{R}_\mathbf{x}$, $\mathcal{U}_\mathbf{x}$, and $\mathcal{C}_\mathbf{x}$ are updated based on the result returned by the chosen verifier v: $\mathcal{R}_\mathbf{x}$ is grown if verification succeeds (v returns 1, cf. Lines 13, 19, and 27); otherwise, $\mathcal{C}_\mathbf{x}$ or $\mathcal{U}_\mathbf{x}$ are grown if counterfactuals are found (v returns -1) or if v times out (v returns 0), respectively (cf. Lines 14 and 30). Finally FaVeX returns $\mathcal{U}_\mathbf{x}$ and $\mathcal{C}_\mathbf{x}$ (cf. Line 31), whose union $\mathcal{U}_\mathbf{x} \cup \mathcal{C}_\mathbf{x}$ yields a v-optimal robust explanation (cf. Definition 3), if the v-opt flag is true, and a robust explanation (cf. Definition 1), otherwise.

▶ **Theorem 5.** *Given a classifier $f : \mathbb{R}^d \to \mathbb{R}^k$, an input vector $\mathbf{x} \in \mathbb{R}^d$, a perturbation radius $\epsilon > 0$, and a verifier $v : Q \to \{-1, 0, 1\}$, FaVeX (Algorithm 3) always terminates and returns disjoint sets $\mathcal{U}_\mathbf{x}, \mathcal{C}_\mathbf{x} \subseteq \{1, \ldots, d\}$ such that $\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}$ is a v-optimal robust explanation of $\mathbf{x}$ (Definition 3), if the v-opt flag is true, or a robust explanations of $\mathbf{x}$ (Definition 1), if the v-opt flag is false.*

Proof. *(Termination)* Algorithm 3 processes a finite set of input feature indexes $\{1, \ldots, d\}$ using a combination of batch splitting and sequential queries. Each step strictly reduces the size of unprocessed batches or classifies a feature index into $\mathcal{R}_\mathbf{x}$, $\mathcal{U}_\mathbf{x}$, or $\mathcal{C}_\mathbf{x}$. Since no feature index is revisited indefinitely and each verifier call terminates (in the worst case due to timeout), the algorithm performs finitely many queries and thus terminates.

(Correctness) Algorithm 3 maintains a partition of the features indexes $\{1, \ldots, d\}$ into invariants $\mathcal{R}_\mathbf{x}$, unknowns $\mathcal{U}_\mathbf{x}$, and counterfactuals $\mathcal{C}_\mathbf{x}$, updated according to the verifier outcome: $v = 1$ adds to $\mathcal{R}_\mathbf{x}$, $v = 0$ to $\mathcal{U}_\mathbf{x}$, and $v = -1$ to $\mathcal{C}_\mathbf{x}$, with robustness queries performed over $\mathcal{R}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}$, when v-opt is true, or over $\mathcal{R}_\mathbf{x}$, v-opt is false. By construction, the resulting partition satisfies the three conditions of Definition 3, when v-opt flag is true, or the condition of Definition 1, when v-opt is false. ◀

Traversal Strategy. The traversal strategy used by FaVeX is shown in Algorithm 4. It leverages an analyzer a (cf. Line 1) to associate a *score* to each input feature index (cf. Line 3). The score is the lower bound of the worst-case logit difference computed by a, when the perturbations are restricted to the currently considered input feature index i (cf. Line 4). The input feature indexes are then sorted in descending order according to their score (cf. Line 5). The underlying intuition is that input feature indexes associated with a

larger lower bound on the worst-case logit difference are more likely to be found invariants for the explanation by FaVeX. Note that our traversal strategy is similar to the one used in VeriX+ [65], with a subtle but important difference: the score associated to each input feature index by VeriX+ [65] is the lower bound on the logit of the true classification output instead of the worst-case logit difference. We argue that prioritizing indexes by their contribution to the worst-case logit difference is more faithful to the verification objective: it directly reflects the margin that must remain positive under perturbations, rather than focusing solely on the robustness of the true class logit without accounting for competing classes. This leads to a traversal strategy that better aligns with the decision criterion used by the verifier and therefore improves the likelihood of identifying invariants early. An empirical comparison of different traversal strategies is shown in Tables 8-11 in Section 7.

5.2 Incremental Branch-and-Bound Verification

We can now remark that most often an invocation of the verifier v within FaVeX concerns a robustness query $(\mathbf{x}, \mathcal{A}, \epsilon, f)$ that differs only slightly from the previous one (typically $\mathcal{A}$ includes only one or, sometimes, few more input feature indexes). Instead of running v from scratch, we propose a strategy that exploits this incremental query evolution to reuse previous verifier computations and, in turn, speed up the overall procedure.

Specifically, we observed that branch-and-bound may perform similar sequences of branching steps for consecutive robustness queries. Thus, we propose to save and reuse the branching steps in later branch-and-bound calls. Concretely, inspired by IVAN [60], we cache the (constraints associated with the) robustness subproblems at the *leaves* of the search tree explored by branch-and-bound, and reuse them as starting point for subsequent branch-and-bound calls. Figure 6 shows an example of branch-and-bound search tree (that occurs during the analysis in Figure 3) where the initial subproblem Q (local robustness verification with 0.6-bounded ℓ_∞ perturbation applied to the input features x_1, x_2, x_3, x_4, x_5, and x_8) is first split into subproblems Q_1 and Q_2 (adding the constraints $\hat{x}_{13} \geq 0$ and $\hat{x}_{13} < 0$, respectively), and then Q_2 is further split into Q_2^1 and Q_2^2 (adding $\hat{x}_{11} \geq 0$ and $\hat{x}_{11} < 0$, respectively). In this case, we cache the (constraints associated with) subproblems Q_1, Q_2^1 and Q_2^2. The subsequent invocation of the verifier (local robustness verification with 0.6-bounded ℓ_∞ perturbation applied to the input features x_1, x_2, x_3, x_4, x_5, x_8, and x_9) will reuse these subproblems (constraints) as a starting point for branch-and-bound, instead of starting from scratch with the robustness query (with an empty set of constraints).

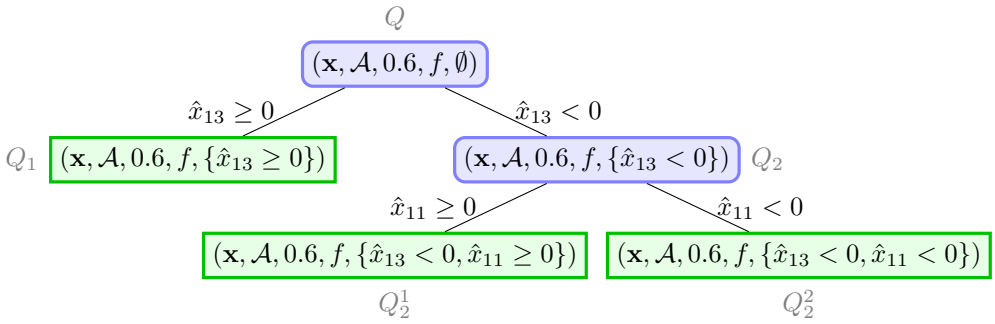

Figure 6 Example of search tree explored by branch-and-bound during the analysis illustrated in Figure 3, where f is the neural network classifier shown in Figure 2, and $\mathcal{A} = \{1, 2, 3, 4, 5, 8\}$.

Algorithm 5 Branch-and-Bound (with Timeout and) Branching Save/Reuse.

```
1: function BAB(a, (x, A, ϵ, f), T, C)        ▷ a : Q̂ → R, (x, A, ϵ, f) ∈ Q, T ∈ N, C ∈ P(P(C))
2:     t1 ← TIME()
3:     for C ∈ C do unresolved ← {(x, A, ϵ, f, C)}                          ▷ Reuse
4:     C ← ∅
5:     for Q ∈ unresolved do
6:         if CEX(x, A, ϵ, f) then
7:             for Q′ ∈ unresolved do C ← C ∪ {Q′_C}                        ▷ Save unresolved
8:             return −1, C
9:         unresolved ← unresolved \ {Q}
10:        if a(Q) > 0 then C ← C ∪ {Q_C}                                   ▷ Save Q
11:        else
12:            t2 ← TIME()
13:            if t2 − t1 < T then
14:                Q1, Q2 ← SPLIT(Q)
15:                unresolved ← unresolved ∪ {Q1, Q2}
16:            else
17:                for Q′ ∈ unresolved do C ← C ∪ {Q′_C}                    ▷ Save unresolved
18:                return 0, C
19:    return 1, C
```

Algorithm 5 shows how we augment branch-and-bound with save and reuse capabilities. When a subproblem Q is verified, we cache the set of constraints Q_C associated with Q (cf. Line 10). Otherwise, if a counterfactual is found (cf. Line 6) or the timeout is hit (cf. Line 16, we cache the constraints of all still unresolved subproblems (cf. Lines 7 and 17). The set of cached constraints is returned together with the verification result (cf. Lines 8, 18, 19) to be reused by subsequent calls to branch-and-bound (cf. Line 3).

In practice, we do not apply this save-and-reuse strategy indiscriminately. While in general it substantially accelerates verification, in some cases, it can be detrimental, e.g., when the number of cached constraints sets becomes too large. To avoid these pitfalls, we enable reuse only under specific conditions and impose (a configurable) limit on the size of the cache. Our practical design choices and heuristics governing when to save and when to reuse constraints of branch-and-bound leaf subproblems are detailed in Section 6.2.

5.3 Restricted-Space Counterfactual Search

The search for counterfactuals during branch-and-bound verification of a robustness query $(\mathbf{x}, \mathcal{A}, \epsilon, f)$ (cf., e.g., Line 4 in Algorithm 1) typically considers the entire set of allowed perturbations $B^{\epsilon}_{\mathcal{A}}(\mathbf{x})$. That is, all input features indexed by $\mathcal{A}$ are treated as freely perturbable, and thus the entire set $\mathcal{A}$ constitutes the counterfactuals *search space*.

However, we can again leverage the dependency between subsequent verification queries in FAVEX (as in Section 5.2), in this case to narrow the counterfactual search space, thereby reducing the search cost. When FAVEX operates in sequential query processing, in particular, input feature indexes are individually and iteratively added to set of feature indexes to which perturbations are allowed. If we have access to a counterfactual $\mathbf{x}' \in B^{\epsilon}_{\mathcal{A}}(\mathbf{x})$ for a robustness query $(\mathbf{x}, \mathcal{A}, \epsilon, f)$, it is extremely likely that x' will be very close to a counterfactual for the next robustness query $(\mathbf{x}, \mathcal{A} \cup B, \epsilon, f)$, in the larger set of allowed perturbations $B^{\epsilon}_{\mathcal{A}\cup B}(\mathbf{x})$.

Figure 7 Illustration of full-space vs. restricted-space counterfactual search. Full-space search (left) explores the entire perturbation region $B^{\epsilon}_{\mathcal{A}\cup B}(\mathbf{x})$. Restricted-space search (right) begins from the output $\mathbf{z}$ of the previous search (which may not be a counterfactual) and varies only the newly added feature(s), thus exploring a much smaller subset of $B^{\epsilon}_{\mathcal{A}\cup B}(\mathbf{x})$.

This motivates our *restricted-space counterfactual search*. Concretely, when solving the next robustness query $(\mathbf{x}, \mathcal{A} \cup B, \epsilon, f)$, instead of conducting the counterfactual search over the entire search space $\mathcal{A} \cup B$ (left of Figure 7), we restrict it to only the newly-added input feature indexes in B (right of Figure 7). At the same time, we center the search space around the result $\mathbf{z} \in B^{\epsilon}_{\mathcal{A}}(\mathbf{x})$ of the previous counterfactual search, which is a point that was heuristically selected to be as close as possible to misclassification. That is, we fix the value of each input feature i indexed by $\mathcal{A}$ to $\mathbf{z}_i$. This way, the counterfactual search explores only a thin slice of $\mathcal{A} \cup B$, but one that is highly promising. In practice, we found this heuristic to be extremely effective for finding valid counterfactuals quickly (cf. Tables 6 and 7 in Section 7). Implementation details are provided in Section 6.2.

6 Implementation

Our implementation[4] is based on the popular deep learning library PyTorch [48].

6.1 Verifier

We adopt the OVAL branch-and-bound framework [6, 4, 11, 10] as the backbone for our verifier, using a fixed timeout of either 300 seconds per query, or 60 seconds, depending on the benchmark. Note that while branch-and-bound is complete in principle, the timeout makes it incomplete in practice (cf. end of Section 3.1), as it is unlikely that the exponential worst-case runtime will be hit during the allotted timeout on non-trivial networks [33].

Bounding Phase. The bounding phase of branch-and-bound requires two components: one tries to prove robustness by operating on a network over-approximation. If the over-approximation is robust, then no counterexampe can exist. This corresponds to computing a lower bound on the worst-case logit difference (cf. Section 3.1). We mainly rely on a popular dual-based algorithm, named α-β-CROWN [63], that solves a dual instance of a network relaxation that replaces the ReLU activations by over-approximating triangles [17]. This algorithm is both employed to bound the logit differences themselves, and to build each network pre-activation (a necessary preliminary step). In line with previous work, we only compute pre-activation bounds once at the branch-and-bound root (i.e., before any

[4] `https://github.com/alessandrodepalma/favex`

splitting) [11]. On smaller networks, we employ a framework configuration that may resort to a tighter network over-approximation for the bounds to the logit differences if needed [10]. We solve up to 2000 branch-and-bound sub-problems in parallel at any time, leveraging GPU acceleration through PyTorch. The other component, based on the evaluation of concrete points, seeks to find a counterexampe violating robustness, upper bounding the worst-case logit difference. In practice, we first run a relatively inexpensive adversarial attack based on a local optimizer, a 10-step Projected Gradient Descent [42] (PGD-10) using the minimal logit difference as loss function, from a single input sampled uniformly within the perturbation region, before entering the branch-and-bound loop. Within branch-and-bound, concrete points to evaluate are collected as a by-product of the over-approximations, and a more expensive local optimizer [14] (500-step MI-FGSM) is repeatedly run.

Branching Phase. In all cases, we use the ReLU splitting partitioning strategy, which operates by splitting ambiguous ReLUs (i.e., their pre-activation can take both negative and positive values for the considered input perturbations) into their two linear phases [5]. In particular, we use a recent strategy [12] that heuristically branches on the ReLU which is deemed to impact the most the tightness of the over-approximation employed during the bounding phase.

Selective MILP Calls. In order to speed up verification on smaller networks, we modified the OVAL framework to allow for the use of a Mixed Integer Linear Programming (MILP) solver whenever a subproblem has fewer than a given number of ambiguous ReLUs. In practice, we use this functionality in two different settings: on very small networks, we run the MILP solver at the root of branch-and-bound (i.e., before any splitting), if the quicker dual-based bounds fail to verify robustness. On harder networks, where this will not pay off, we call the MILP solver when no ambiguous neurons are left, which is in practice faster than ensuring convergence of dual-based bounding on the underlying linear program. MILPs are solved using Gurobi, a commercial black-box solver [21].

6.2 FaVeX

Restricted-Space Counterfactual Search. We run the novel restricted-space counterfactual search described in Section 5.3 only if the preliminary PGD-10 attack (see Section 6.1) fails to find counterexamples. We execute several (i.e., 128) searches for counterexamples in parallel over the GPU, starting from different points in the restricted search space. Specifically, we include both the extrema of the search-space (the interval lower and upper bounds), and a series of starting points sampled uniformly within the interval. PGD-10 [42] is then run for each of these starting points, using the minimal logit difference as loss function: we call this implementation **Restricted-Space Attack (RSA)**.

Traversal. Our traversal strategy (cf. Section 5.1) requires as many calls to the analyzer a as the number of input features. In practice, we execute a number of these in parallel on a GPU. We consider two analyzers: α-CROWN [67], and the less expensive IBP [19, 18].

Leaf Reuse. We implemented the leaf reuse described in Section 5.2 through a direct modification of the OVAL branch-and-bound code. We store leaves as a list of branching decisions, which are then enforced after the computation of the pre-activation bounds at the root. In practice, we do not necessarily inherit the leaves from the previous branch-and-bound

Table 1 Size of the network architectures employed in the experimental evaluation. For convolutional networks, the number of activations is computed for GTSRB and CIFAR-10 inputs.

Network	N. layers	N. activations	N. parameters
`FC-10x2`	3	2.00×10^1	3.10×10^4
`FC-50x2`	3	1.00×10^2	1.57×10^5
`CNN-3`	3	2.46×10^4	2.18×10^5
`CNN-7`	7	2.30×10^5	1.72×10^7

call. For efficiency reasons, we may discard leaves (i.e., start branch-and-bound from scratch) when too many of them are stored (more than a tunable threshold). Analogously, we discard leaves in case the previous call is a timeout and FaVeX is processing robustness queries sequentially (cf. Section 5), as inheriting them would likely cause further timeouts. Finally, in case of timeouts during the binary search batch processing mode, we inherit the leaves from the branch-and-bound call preceding the timeout.

7 Experimental Evaluation

This section presents an experimental evaluation of verifier-optimal robust explanations (Section 7.2), of FaVeX (Section 7.3), and of the associated traversal strategies (Section 7.4), preceded by a description of the experimental setting (Section 7.1).

7.1 Experimental Setting

We carry out experiments on computer vision datasets popular in the formal explanations literature [66, 65]. In particular, we consider MNIST [36], a black-and-white handwritten digit classification task; a 10-class version of GTSRB [58] taken from previous work [66], a traffic sign recognition task; and CIFAR-10 [34], a popular 10-class image classification benchmarks (examples of classes including “airplane”, “automobile”, “cat”, “dog”).

Network Architectures. In order to assess the scalability of FaVeX and of the presented definition, we consider 4 different feedforward ReLU networks of varying sizes, detailed in Table 1. The hardness of neural network verification mainly depends on the number of activations acting non-linearly. `FC-10x2` is a fully connected network with two hidden layers of width equal to 10. `FC-50x2` displays the same structure, but with hidden layers of width 50. Both networks have relatively few activations, and were trained using standard (SGD-based) network training and ℓ_1 regularization. `CNN-7` is a 7-layer convolutional network that is commonly employed in the certified training literature (networks trained for verified robustness), where it reaches state-of-the-art performance [12]: we use this to showcase scalability on networks relevant to the broader robust machine learning literature, and take the trained networks directly from previous work [12]. As `CNN-7` has 2.3×10^5 activations, we also train (using the same training scheme) a narrower 3-layer version, named `CNN-3`, which has 2.46×10^4 activations. `CNN-7` and `CNN-3` were trained to be robust against perturbations of $\epsilon = 0.2$ and $\epsilon = \frac{2}{255}$ on MNIST, and CIFAR-10, respectively.

Computational Setup. We run the experimental evaluation on a single GPU, using 6 CPU cores and allocating 20GB of RAM. In order to run multiple experiments at once, we use different machines, but allocate each combination of network and dataset to a specific machine for consistency. We employed the following GPU models: Nvidia RTX 6000, Nvidia RTX 8000. All machines are equipped with the following CPU: AMD EPIC 7302.

Tuning. Before launching the experiments, we tested different configurations of the OVAL framework in order to minimize the overall runtime. We found it beneficial to avoid the custom branch-and-bound, and to call the MILP solver before any branching (see Section 6.1) on all the `FC-10x2` experiments, and on the `FC-50x2` MNIST experiments. The only hyper-parameter associated to FaVeX is k, the maximal number of leaves to be stored for reuse before discarding them. We tune $k \in \{0, 25, 100, 500, 5000, 50000\}$ on test images not employed for the evaluation. We leave the number of leaves unbounded if the tuning points to $k = 50000$: this is the case on the fully-connected network instances for which we do not directly resort to the MILP solver (`FC-50x2` on GTSRB and CIFAR-10).

7.2 Scalable Explanations

We here evaluate the scalability improvements associated to the definition of verifier-optimal robust explanations, presented in Section 4. Specifically, we use FaVeX and a fixed traversal order (α-FaVeX for CIFAR-10 and FaVeX-IBP for MNIST) to compute both (standard) robust explanations and `OVAL`-optimal robust explanations, using a per-query timeout of 60 seconds on `CNN-3` and `CNN-7`. Tables 2-3 and Figure 8 show that `OVAL`-optimal robust explanations are marginally (less than 3%) larger than standard robust explanations. On the other hand, standard robust explanations are extremely expensive to compute, ranging from roughly 45 minutes per image on MNIST for `CNN-3`, to above 7 hours per image on CIFAR-10 for `CNN-7`. Note that, despite this computational effort, no counterexamples can be found (cf. end of Section 3.2). On the other hand, `OVAL`-optimal robust explanations are significantly faster to compute, and result in a significant number of counterfactuals, clearly showing the superior scalability of the proposed definition.

7.3 Computing Explanations Faster

We now turn to evaluating the efficacy of FaVeX when computing both standard robust explanations, and verifier-optimal robust explanations, again using a fixed traversal strategy

Table 2 On `CNN-3`, `OVAL`-optimal robust explanations are significantly more scalable than standard robust explanations computed using the same verifier. Results are averaged over the first 10 images of the test set. Entries highlighted in bold are the smallest average runtime, and the largest average in number of found counterfactuals.

	MNIST, $\epsilon = 0.25$				CIFAR-10, $\epsilon = \frac{8}{255}$			
Robust Explanation	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]
Standard	247.80	0.00	247.80	2689.77	461.00	0.00	461.00	8984.10
V-optimal	254.70	**160.30**	94.40	**612.75**	462.10	**210.40**	251.70	**1150.82**

Table 3 `OVAL`-optimal robust explanations are significantly more scalable than standard robust explanations computed using the same verifier on `CNN-7`. We average results over the first 3 images of the test set, and highlight in bold the smallest average runtime the largest average in number of found counterfactuals.

	MNIST, $\epsilon = 0.25$				CIFAR-10, $\epsilon = \frac{16}{255}$			
Robust Explanation	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]
Standard	452.00	0.00	452.00	14 317.57	730.67	0.00	730.67	25 487.29
V-optimal	456.66	**207.33**	249.33	**4446.53**	733.33	**467.00**	266.33	**6532.98**

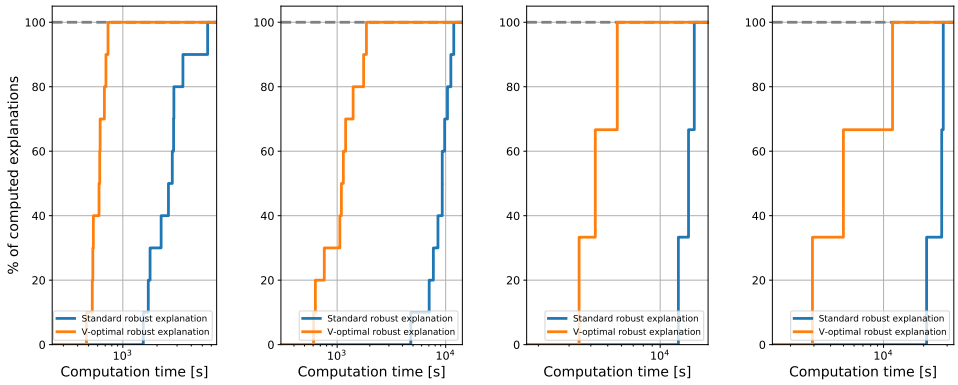

(a) `CNN-3`, MNIST. **(b)** `CNN-3`, CIFAR-10. **(c)** `CNN-7`, MNIST. **(d)** `CNN-7`, CIFAR-10.

Figure 8 Percentage of computed explanations as a function of runtime for Tables 2 and 3.

(α-FaVeX for CIFAR-10 and GTSRB, FaVeX-IBP for MNIST). All considered algorithms to compute the explanations rely on the OVAL framework as verifier, using the same configuration (see Section 6.1). We compare against the following baselines:

- Sequential, which processes all the input features sequentially (cf. Section 5), mirroring what done in VeriX [66].
- Binary Search, which exclusively uses the binary search-based batch processing, mirroring VeriX+ [65].

Differently from the original works [66, 65], we use the OVAL verifier in both cases for consistency. In order to assess the effect of each of the components of FaVeX in isolation, we evaluate the following methods, which are all novel in the context of verified explanations:

- BinS + Incr, which uses incremental branch-and-bound (leaf reuse) as described in Section 5.2 on top of Binary Search.
- BinS + Incr + RSA, which employs the Restricted Space Attack (RSA) implementation of the counterfactual search presented in Section 5.3 on top of BinS + Incr.
- FaVeX, which features the three following additions compared to Binary Search: (i) leaf reuse, (ii) RSA, and (iii) the fallback to sequential processing presented in Section 5.1.

Standard Robust Explanations. Tables 4-5 and Figure 9 show that FaVeX is significantly faster than both baselines on all considered settings on both `FC-10x2` and `FC-50x2`. Binary Search is consistently better than sequentially considering the input features on all benchmarks, cutting runtime almost by a third on MNIST for `FC-10x2`. FaVeX, however, reduces runtime even further, attaining speedup factors of up to 13× on `FC-10x2` and up to 16× on `FC-50x2`. Note that on `FC-10x2` and for MNIST on `FC-50x2`, where the MILP solver is called at the branch-and-bound root, no explicit OVAL branching is carried out (see Section 6.1), so leaf reuse is not employed (hence the / markers for BinS + Incr in Table 5 and its absence in Table 4). As visible from the performance of BinS + Incr + RSA, most of the remarkable improvements in these settings are to be attributed to the reduced-space counterfactual search. The fallback to sequential processing provides an additional speedup on most of the benchmarks, for instance cutting average runtime by more than 2× on GTSRB for `FC-10x2`. Incremental branch-and-bound (BinS + Incr) also achieves significant speed-ups on `FC-50x2`: more than a factor 1.7× on both GTSRB and CIFAR10, pointing to its efficacy on computing standard robust explanations for small networks.

Table 4 FaVeX attains significant speed-ups against previous algorithms to compute standard robust explanations [66, 65] on `FC-10x2`. Results are averaged over the first 100 images of the test set. Bold entries correspond to the smallest average runtime.

	MNIST, $\epsilon = 0.1$		GTSRB, $\epsilon = 0.05$	
Method	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]
Sequential	70.90	182.62	352.26	518.34
Binary Search	70.90	65.21	352.26	345.48
BinS + Incr + RSA	70.90	25.98	352.26	56.26
FaVeX	70.90	**21.21**	352.26	**26.08**

Table 5 FaVeX is significantly faster than previous algorithms to compute standard robust explanations [66, 65] on `FC-50x2`. Results are averaged over the first 100 images of the test set. Bold entries correspond to the smallest average runtime.

	MNIST, $\epsilon = 0.2$		GTSRB, $\epsilon = 0.1$		CIFAR-10, $\epsilon = \frac{2}{255}$	
Method	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]
Sequential	82.74	243.38	287.02	688.24	204.23	468.23
Binary Search	82.74	146.45	287.02	476.55	204.23	197.55
BinS + Incr	/	/	286.99	275.47	204.23	110.91
BinS + Incr + RSA	82.74	58.33	287.01	55.40	204.23	25.13
FaVeX	82.74	**48.16**	287.02	**30.48**	204.23	**13.92**

Verifier-Optimal Robust Explanations. As visible from Tables 6-7 and Figure 10, FaVeX is also beneficial when computing `OVAL`-optimal robust explanations on larger networks, with speed-ups up to 67% and 79% on MNIST for `CNN-3` and for `CNN-7`, respectively. Remarkably, on MNIST, the use of FaVeX also results in a larger number of counterfactuals being found, pointing to its superior efficacy in terms of counterfactual search: as visible by comparing BinS + Incr + RSA with BinS + Incr, this is due to the efficacy of the reduced-space attack. Interestingly, Sequential is almost always preferable to Binary Search on these

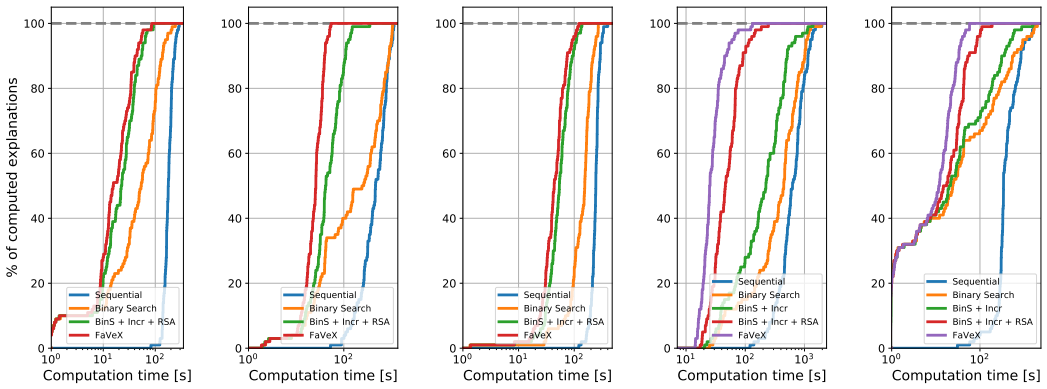

(a) `FC-10x2`, MNIST. **(b)** `FC-10x2`, GTSRB. **(c)** `FC-50x2`, MNIST. **(d)** `FC-50x2`, GTSRB. **(e)** `FC-50x2`, CIFAR10.

Figure 9 Percentage of computed explanations as a function of runtime for Tables 4 and 5.

Table 6 On `CNN-3`, FaVeX computes `OVAL`-optimal robust explanations faster than previous computation strategies [66, 65] while at the same time finding more counter-factuals. Results are averaged over the first 10 images of the test set. Bold entries correspond to the smallest average runtime and the largest average in number of provided counterfactuals.

	MNIST, $\epsilon = 0.25$			CIFAR-10, $\epsilon = \frac{8}{255}$		
Method	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]
Sequential	129.10	125.50	1164.89	209.30	253.10	1287.96
Binary Search	134.20	120.40	1026.67	209.60	252.70	1367.25
BinS + Incr	135.20	119.50	1040.82	209.60	252.50	1374.40
BinS + Incr + RSA	157.70	96.60	727.94	**211.90**	250.20	1334.00
FaVeX	**160.30**	94.40	**612.75**	210.40	251.70	**1150.82**

Table 7 On `CNN-7`, FaVeX computes `OVAL`-optimal robust explanations faster than previous computation strategies [66, 65] while at the same time finding more counter-factuals. Results are averaged over the first 3 images of the test set. Bold entries correspond to the smallest average runtime and the largest average in number of provided counterfactuals.

	MNIST, $\epsilon = 0.25$			CIFAR-10, $\epsilon = \frac{16}{255}$		
Method	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]
Sequential	142.33	315.00	7979.61	464.00	269.33	6868.53
Binary Search	148.33	308.67	9165.06	464.33	269.00	8284.12
BinS + Incr	151.33	305.00	9116.13	467.00	266.33	8231.01
BinS + Incr + RSA	196.33	260.33	5406.70	**468.00**	265.33	8304.52
FaVeX	**207.33**	249.33	**4446.53**	467.00	266.33	**6532.98**

networks: this is due to the relatively large size of the explanations. In fact, for larger explanations, Binary Search will go relatively deep in most branches of its recursion tree, hence resulting in a larger number of calls to the verifier. This is reflected in the relative benefit of the fallback to sequential processing (visible by comparing FaVeX with BinS + Incr + RSA), which is the only FaVeX component yielding consistent speed-ups on all considered `CNN-3` and `CNN-7` benchmarks. Incremental branch-and-bound does not appear to be beneficial on these larger networks: we ascribe this to the significantly larger activation space, which makes it unlikely for a branching decision to be effective across queries to the verifier. Finally, it is worth pointing out that the time to compute the explanations sharply increases with the size of the network. Focusing on MNIST, this takes FaVeX 21.21 seconds on average for `FC-10x2`, 48.16 on `FC-50x2`, 612.75 seconds on `CNN-3`, and 4446.53 seconds on `CNN-3`.

7.4 Traversal Strategies Analysis

We now turn our attention to studying the effect of the employed traversal strategies, comparing α-FaVeX and FaVeX-IBP with the VeriX [66] and VeriX+ [65] traversals. We use FaVeX to compute the explanations owing to its superior speed.

Standard Robust Explanations. Tables 8 and 9 show the results of the analysis for standard robust explanations on smaller networks (`FC-10x2` and `FC-50x2`). While α-FaVeX consistently yields the smallest explanations on GTSRB and CIFAR-10, it underperforms

Table 8 Comparison of different traversal strategies for standard robust explanations on `FC-10x2`. Results are averaged over the first 100 images of the test set. Bold entries correspond to the smallest average explanation size.

	MNIST, $\epsilon = 0.1$		GTSRB, $\epsilon = 0.05$	
Traversal	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]
VeriX	91.26	38.10	737.02	37.49
VeriX+	72.08	23.05	357.26	24.06
α-FaVeX	79.85	30.24	**352.26**	26.08
FaVeX-IBP	**70.90**	21.21	355.52	24.32

Table 9 Comparison of different traversal strategies for standard robust explanations on `FC-50x2`. Results are averaged over the first 100 images of the test set. Bold entries correspond to the smallest average explanation size.

	MNIST, $\epsilon = 0.2$		GTSRB, $\epsilon = 0.1$		CIFAR-10, $\epsilon = \frac{2}{255}$	
Traversal	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]
VeriX	94.02	85.00	851.64	117.20	310.56	18.40
VeriX+	**77.46**	48.17	289.56	25.74	222.66	13.85
α-FaVeX	102.35	67.01	**287.02**	30.48	**204.23**	13.92
FaVeX-IBP	82.74	48.16	287.25	25.31	220.96	14.02

on MNIST, where strategies based on IBP attain superior performance. Nevertheless, on GTSRB and CIFAR-10, except for the VeriX strategy, which appears to produce markedly larger explanations, the performance of the considered strategies is not particularly dissimilar, with the best size reduction of α-FaVeX (attained on CIFAR-10 for `FC-50x2`) compared to VeriX+ being 9%. Finally, we remark that `FC-50x2` appears to be more robust than the smaller `FC-10x2`: on GTSRB, smaller explanations are produced by all traversals except VeriX, in spite of the twice-larger considered perturbation radius.

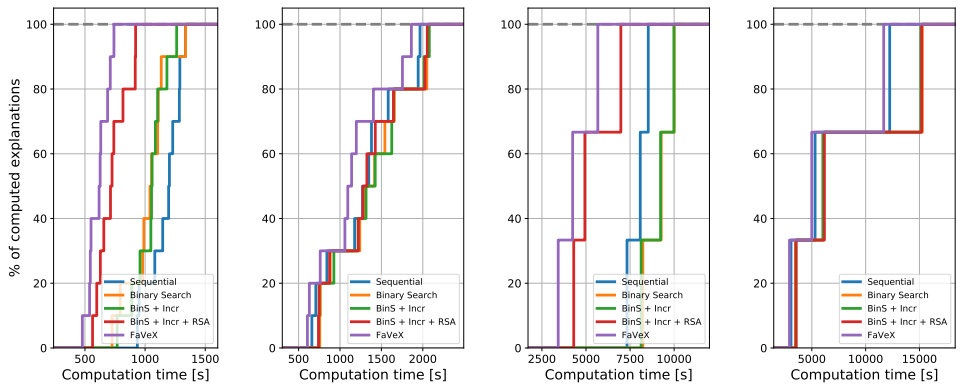

(a) `CNN-3`, MNIST. **(b)** `CNN-3`, CIFAR-10. **(c)** `CNN-7`, MNIST. **(d)** `CNN-7`, CIFAR-10.

Figure 10 Percentage of computed explanations as a function of runtime for Tables 6 and 7.

Verifier-Optimal Robust Explanations. Tables 10 and 11 study traversals for OVAL-optimal robust explanations on the larger `CNN-3` and `CNN-7`. In order to allow for a coherent comparison, we here consider the explanation as the union of $\mathcal{C}_\mathbf{x}$ and $\mathcal{U}_\mathbf{x}$. As for the smaller networks, VeriX+ produces smaller explanations on MNIST. On the other hand, the margins of improvement of α-FaVeX are reduced, with FaVeX-IBP producing marginally smaller explanations on CIFAR-10 for `CNN-7`. At the same time, α-FaVeX non-negligibly reduces the number of counterfactuals on CIFAR-10 for both networks, highlighting that the allocation of pixels between $\mathcal{C}_\mathbf{x}$ and $\mathcal{U}_\mathbf{x}$ depends on the traversal strategy.

Table 10 Comparison of different traversal strategies for `OVAL`-optimal robust explanations on `CNN-3`. Results are averaged over the first 10 images of the test set. Bold entries correspond to the smallest $\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}$.

	MNIST, $\epsilon = 0.25$				CIFAR-10, $\epsilon = \frac{8}{255}$			
Method	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]
VeriX	282.80	160.50	122.30	974.35	765.50	437.20	328.30	1925.09
VeriX+	**247.40**	155.20	92.20	576.67	468.90	262.00	206.90	915.45
α-FaVeX	289.90	181.20	108.70	696.31	**462.10**	210.40	251.70	1150.82
FaVeX-IBP	254.70	160.30	94.40	612.75	466.40	250.90	215.50	941.95

Table 11 Comparison of different traversal strategies for `OVAL`-optimal robust explanations on `CNN-7`. Results are averaged over the first 3 images of the test set. Bold entries correspond to the smallest $\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}$.

	MNIST, $\epsilon = 0.25$				CIFAR-10, $\epsilon = \frac{16}{255}$			
Method	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]	$\lvert\mathcal{C}_\mathbf{x} \cup \mathcal{U}_\mathbf{x}\rvert$	$\lvert\mathcal{C}_\mathbf{x}\rvert$	$\lvert\mathcal{U}_\mathbf{x}\rvert$	time [s]
VeriX	547.00	123.33	423.67	7763.91	945.00	728.67	216.33	4974.69
VeriX+	**429.33**	196.67	232.67	4394.67	735.67	522.00	213.67	5105.86
α-FaVeX	552.33	234.67	317.67	5344.59	733.33	467.00	266.33	6532.98
FaVeX-IBP	456.67	207.33	249.33	4446.53	**729.33**	512.67	216.67	5099.98

7.5 Ablation: Cap on Stored Leaves

FaVeX features only a single hyper-parameter: k, the maximum number of leaves to be stored for reuse, tuned as described in Section 7.1. As outlined in Section 6.2, when the number of branch-and-bound leaves from a call to the verifier is greater than k, branch-and-bound for the next verifier call will start from scratch, discarding any leaf information. In order to report the effect of k on final performance, we here show results for three values: one relatively small ($k = 25$), one intermediate ($k = 500$), and one associated to uncapped storage ($k = \infty$). Consistently with our focus on the efficiency of FaVeX, we report results for the setups from Section 7.3, excluding benchmarks where the MILP solver is called at the branch-and-bound root. Tables 12, 13 and 14 show that FaVeX is relatively insensitive to its hyper-parameter, and that the tuning process consistently leads to good performance. Consistently with the fact that leaf reuse was found to be ineffective on `CNN-3` and `CNN-7` in Section 7.3, allowing for the reuse of a large number of leaves leads has a negative impact on performance in Tables 13 and 14.

Table 12 Ablation on k, the maximum number of leaves that can be stored for reuse within FaVeX, for the experiments from Table 5. Results are averaged over the first 100 images of the test set. Bold entries correspond to the tuned k value employed in the original experiment.

	GTSRB, $\epsilon = 0.1$		CIFAR-10, $\epsilon = \frac{2}{255}$	
k	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{E}_{\mathbf{x}}\rvert$	time [s]
25	287.02	40.95	204.23	14.44
500	287.02	29.42	204.23	13.90
∞	**287.02**	**30.48**	**204.23**	**13.92**

Table 13 Ablation for the maximum number of leaves that can be stored for reuse within FaVeX (denoted k) for the experiments from Table 6. Results are averaged over the first 10 images of the test set. Bold entries correspond to the tuned k value employed in the original experiment.

	MNIST, $\epsilon = 0.25$			CIFAR-10, $\epsilon = \frac{8}{255}$		
k	$\lvert\mathcal{C}_{\mathbf{x}}\rvert$	$\lvert\mathcal{U}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{C}_{\mathbf{x}}\rvert$	$\lvert\mathcal{U}_{\mathbf{x}}\rvert$	time [s]
25	**160.30**	**94.40**	**612.75**	210.00	252.20	1139.99
500	160.20	94.20	629.48	**210.40**	**251.70**	**1150.82**
∞	159.00	96.00	785.96	209.80	253.10	1204.79

Table 14 Ablation on k, the maximum number of leaves that can be stored for reuse within FaVeX, for the experiments from Table 7. Results are averaged over the first 3 images of the test set. Bold entries correspond to the tuned k value employed in the original experiment.

	MNIST, $\epsilon = 0.25$			CIFAR-10, $\epsilon = \frac{16}{255}$		
k	$\lvert\mathcal{C}_{\mathbf{x}}\rvert$	$\lvert\mathcal{U}_{\mathbf{x}}\rvert$	time [s]	$\lvert\mathcal{C}_{\mathbf{x}}\rvert$	$\lvert\mathcal{U}_{\mathbf{x}}\rvert$	time [s]
25	**207.33**	**249.33**	**4446.53**	466.00	267.33	6542.63
500	186.33	270.00	4889.95	**467.00**	**266.33**	**6532.98**
∞	207.33	252.67	4658.70	463.67	272.67	6759.26

8 Related Work

A large body of work from the machine learning community has sought to provide explanations on the decisions of machine learning models. Earlier works from the machine learning community, for instance, propose to build interpretable approximations of the original model [41, 51, 49]. Another line of work tries to find points close to the input such that the decision would change (counterfactuals), which are robust in either the input space [37], or the parameter space [22, 30]. For a survey on robust counterfactual explanation in machine learning, we direct the reader to [31]. At the same time, multiple approaches to heuristically select a subset of features deemed responsible for the predictions exist, typically based on model gradients [54, 52]. However, these heuristics lack any type of formal guarantees, which are however crucial for safety-critical systems. Earlier work on formal explainability was carried out by the logic community, focusing on abductive explanations [45, 43, 46] (cf. Section 3.2). This was followed by work that adds locality constraints, necessary to obtain meaningful explanations on complex models such as neural networks, first in natural language processing [35], and then in computer vision, where the computation of formal explanations has been defined through a series of robustness queries to a neural network verifier [66, 65].

Earlier approaches to neural network verification typically relied on the use of black-box solvers, either from the model checking community [33], or from mathematical optimization [59, 39]. All these complete approaches can be seen as specific instances of branch-and-bound [6]. Incomplete approaches, instead, rely on network over-approximations, and were derived in parallel by both the formal methods [38, 18, 57, 56, 62, 47], and the machine learning [19, 69] communities, often relying on tools from optimization [64, 16]. A series of incomplete verifiers has been designed and integrated within branch-and-bound (as bounding algorithm), yielding verifiers that scale significantly better than earlier approaches based on black-box solvers [4, 67, 12, 25]. The current state-of-the-art, as highlighted by yearly competitions in the area [3], relies on fast incomplete algorithms based on linear bound propagation techniques [63], possibly with the addition of optimizations such as custom cutting planes [68, 71]. A recent work [60], IVAN, proposes to reuse information from previous branch-and-bound calls when verifying slightly modified versions of the same network. Our incremental approach, presented in Section 5.2 is a simplified version of this idea yet applied to formal explainability, where the network stays the same, but the verifier is called on a series of similar input domains that differ in the perturbed features of the same input point.

9 Conclusion and Future Work

While verified explanations offer a theoretically-principled approach for explainable machine learning, their applicability is severely limited by their scalability. This work makes a step towards improving the scalability of verified explanations for neural networks classifiers by both introducing a novel algorithm for their computation, and presenting a novel and more realistic definition of verified explanations.

We proposed FaVeX (Section 5), a new algorithm that substantially accelerates the computation of verified explanations. Compared to previous work, FaVeX presents three key improvements:

- **(i)** it dynamically combines both batch processing of robustness queries in a binary-search fashion and their sequential processing, getting the best of both worlds depending on the query (Section 5.1);
- **(ii)** it reuses information from previous executions of branch-and-bound, thus accelerating the overall verification process (Section 5.2);
- **(iii)** it leverages a novel reduced-space counterfactual search that prevents calls to branch-and-bound by effectively re-using information from previous queries, which are used to quickly find counterfactuals (Section 5.3).

Furthermore, we introduced *verifier-optimal robust explanations* (Definition 3), which explicitly account for the practical incompleteness of neural network verifiers on medium-sized networks by allowing features whose robustness cannot be practically ascertained to be perturbed along the invariants, and produce a hierarchical explanation. In the case of a perfect verifier, the presented definition reverts to the standard definition of an optimal robust explanation from previous work (Lemma 4).

Our experiments demonstrate that FaVeX yields speed-ups of up to an order of magnitude compared to previous algorithms when computing standard robust explanations on small fully-connected networks, and that it cuts explanation times while discovering more counterfactuals when computing verifier-optimal robust explanations on larger convolutional networks. At the same time we show that, using FaVeX, verifier-optimal robust explanations can be computed for a fraction of the cost of the standard definition of robust explanations, while at the same time allowing for the production of several counterfactuals, which cannot be

otherwise provided. In particular, FAVEX can find hundreds of counterfactuals on state-of-the-art networks trained for provable robustness to small adversarial examples, which feature hundreds of thousands of neurons, hence producing formal explanations at scale.

Threats to Validity. While our experimental evaluation demonstrates improvements over prior approaches, several factors may affect the generalizability of our findings:

Network Architectures. Our experiments inherit the current limitations of state-of-the-art neural network verification. We focus on piecewise-linear feedforward neural networks with ReLU activations, which are more amenable to verification in practice. Further advances in verification techniques are a prerequisite for scalably extending verified explanations to other architectures.

Computational Resources. All experiments were conducted with fixed computational resources (6 CPU cores, 20GB RAM, single GPU). While different hardware configurations or resource allocations may affect absolute runtimes, we do not expect them to significantly alter the observed trends.

Domain Specificity. Consistently with previous work, our evaluation focuses on computer vision tasks (MNIST, GTSRB, CIFAR-10). The effectiveness of FAVEX and verifier-optimal robust explanations on other domains remains to be investigated. As the proposed algorithm is not inherently tied to image data, we do not expect fundamentally different behavior on other datasets (the example in Section 2 originates from a tabular dataset). On the other hand, domain-specific challenges (for instance, the discreteness of natural language processing data) may require substantial adaptations to verification and formal explainability techniques.

Verifier Choice. Our results are specific to the OVAL branch-and-bound framework with α-β-CROWN bounding. While different verifiers or verification algorithms may exhibit variations in absolute performance, we expect the qualitative trends observed for FAVEX to generalize across verification backends.

Future Work. While this work substantially improves on the scalability of formal explanations, their computation still requires in the order of an hour per image on networks with tens of millions of parameters and trained for verifiability. These are small compared to real-world applications in many domains. Further progress on scalability is hence fundamental to ensure their widest applicability. Promising directions in this sense include: exploring more complex ways to re-use information from the various queries to the verifier, systematically exploring the interaction between training methods and formal explanability, and devising ad-hoc algorithms to train networks that combine good performance and formal explanability. Finally, while our experiments focus on vision models, testing and potentially adapting verifier-optimal robust explanations and FAVEX to other data domains is an exciting avenue for future work.

References

1 Aws Albarghouthi. Introduction to neural network verification. *CoRR*, abs/2109.10317, 2021. arXiv:2109.10317.

2 Ross Anderson, Joey Huchette, Will Ma, Christian Tjandraatmadja, and Juan Pablo Vielma. Strong mixed-integer programming formulations for trained neural networks. *Math. Program.*, 183(1):3–39, 2020. doi:10.1007/S10107-020-01474-5.

3 Christopher Brix, Mark Niklas Müller, Stanley Bak, Taylor T Johnson, and Changliu Liu. First three years of the international verification of neural networks competition (vnn-comp). *International Journal on Software Tools for Technology Transfer*, 25(3):329–339, 2023. doi:10.1007/S10009-023-00703-4.

4 Rudy Bunel, Alessandro De Palma, Alban Desmaison, Krishnamurthy Dvijotham, Pushmeet Kohli, Philip HS Torr, and M Pawan Kumar. Lagrangian decomposition for neural network verification. In *Conference on Uncertainty in Artificial Intelligence*, 2020.

5 Rudy Bunel, Jingyue Lu, Ilker Turkaslan, Philip H. S. Torr, Pushmeet Kohli, and M. Pawan Kumar. Branch and bound for piecewise linear neural network verification. *J. Mach. Learn. Res.*, 21:42:1–42:39, 2020. URL: https://jmlr.org/papers/v21/19-468.html.

6 Rudy Bunel, Ilker Turkaslan, Philip HS Torr, Pushmeet Kohli, and M Pawan Kumar. A unified view of piecewise linear neural network verification. In *Neural Information Processing Systems*, 2018.

7 John W. Chinneck and Erik W. Dravnieks. Locating minimal infeasible constraint sets in linear programs. *INFORMS J. Comput.*, 3(2):157–168, 1991. doi:10.1287/IJOC.3.2.157.

8 Adnan Darwiche and Auguste Hirth. On the reasons behind decisions. In Giuseppe De Giacomo, Alejandro Catalá, Bistra Dilkina, Michela Milano, Senén Barro, Alberto Bugarín, and Jérôme Lang, editors, *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 712–720. IOS Press, 2020. doi:10.3233/FAIA200158.

9 Adnan Darwiche and Auguste Hirth. On the (complete) reasons behind decisions. *J. Log. Lang. Inf.*, 32(1):63–88, 2023. doi:10.1007/S10849-022-09377-8.

10 Alessandro De Palma, Harkirat Singh Behl, Rudy Bunel, Philip H. S. Torr, and M. Pawan Kumar. Scaling the convex barrier with sparse dual algorithms. *Journal of Machine Learning Research*, 2024.

11 Alessandro De Palma, Rudy Bunel, Alban Desmaison, Krishnamurthy Dvijotham, Pushmeet Kohli, Philip H. S. Torr, and M. Pawan Kumar. Improved branch and bound for neural network verification via lagrangian decomposition. *CoRR*, abs/2104.06718, 2021. arXiv:2104.06718.

12 Alessandro De Palma, Rudy Bunel, Krishnamurthy (Dj) Dvijotham, M. Pawan Kumar, Robert Stanforth, and Alessio Lomuscio. Expressive losses for verified robustness via convex combinations. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL: https://openreview.net/forum?id=mzyZ4wzKlM.

13 Alessandro De Palma, Greta Dolcetti, and Caterina Urban. FaVeX. Software (visited on 2026-06-09). URL: https://github.com/alessandrodepalma/favex, doi:10.4230/artifacts.26386.

14 Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9185–9193, 2018. doi:10.1109/CVPR.2018.00957.

15 Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint*, 2017. arXiv:1702.08608.

16 Krishnamurthy Dvijotham, Robert Stanforth, Sven Gowal, Timothy Mann, and Pushmeet Kohli. A dual approach to scalable verification of deep networks. In *Conference on Uncertainty in Artificial Intelligence*, 2018.

17 Ruediger Ehlers. Formal verification of piece-wise linear feed-forward neural networks. *Automated Technology for Verification and Analysis*, 2017.

18 Timon Gehr, Matthew Mirman, Dana Drachsler-Cohen, Petar Tsankov, Swarat Chaudhuri, and Martin T. Vechev. AI2: safety and robustness certification of neural networks with abstract interpretation. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings,*

21-23 May 2018, San Francisco, California, USA, pages 3–18. IEEE Computer Society, 2018. doi:10.1109/SP.2018.00058.

19 Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Arthur Mann, and Pushmeet Kohli. Scalable verified training for provably robust image classification. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 4841–4850. IEEE, 2019. doi:10.1109/ICCV.2019.00494.

20 Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):93:1–93:42, 2019. doi:10.1145/3236009.

21 Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual. URL: https://www.gurobi.com.

22 Faisal Hamman, Erfaun Noorani, Saumitra Mishra, Daniele Magazzeni, and Sanghamitra Dutta. Robust counterfactual explanations for neural networks with probabilistic guarantees. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 12351–12367. PMLR, 23–29 July 2023. URL: https://proceedings.mlr.press/v202/hamman23a.html.

23 Yotam Hechtlinger. Interpretation of Prediction Models Using the Input Gradient. *CoRR arXiv*, 2016. arXiv:1611.07634.

24 Fred Hemery, Christophe Lecoutre, Lakhdar Sais, and Frédéric Boussemart. Extracting mucs from constraint networks. In Gerhard Brewka, Silvia Coradeschi, Anna Perini, and Paolo Traverso, editors, *ECAI 2006, 17th European Conference on Artificial Intelligence, August 29 - September 1, 2006, Riva del Garda, Italy, Including Prestigious Applications of Intelligent Systems (PAIS 2006), Proceedings*, volume 141 of *Frontiers in Artificial Intelligence and Applications*, pages 113–117. IOS Press, 2006. URL: http://www.booksonline.iospress.nl/Content/View.aspx?piid=1657.

25 Patrick Henriksen and Alessio Lomuscio. DEEPSPLIT: an efficient splitting method for neural network verification via indirect effect analysis. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 2549–2555. ijcai.org, 2021. doi:10.24963/IJCAI.2021/351.

26 Xuanxiang Huang and João Marques-Silva. On the failings of shapley values for explainability. *Int. J. Approx. Reason.*, 171:109112, 2024. doi:10.1016/J.IJAR.2023.109112.

27 Alexey Ignatiev, Nina Narodytska, and João Marques-Silva. Abduction-based explanations for machine learning models. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 1511–1519. AAAI Press, 2019. doi:10.1609/AAAI.V33I01.33011511.

28 Alexey Ignatiev, Nina Narodytska, and João Marques-Silva. On relating explanations and adversarial examples. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 15857–15867, 2019. URL: https://proceedings.neurips.cc/paper/2019/hash/7392ea4ca76ad2fb4c9c3b6a5c6e31e3-Abstract.html.

29 Yacine Izza, Xuanxiang Huang, António Morgado, Jordi Planes, Alexey Ignatiev, and João Marques-Silva. Distance-restricted explanations: Theoretical underpinnings & efficient implementation. In Pierre Marquis, Magdalena Ortiz, and Maurice Pagnucco, editors, *Proceedings*

of the 21st International Conference on Principles of Knowledge Representation and Reasoning, KR 2024, Hanoi, Vietnam. November 2-8, 2024, 2024. doi:10.24963/KR.2024/45.

30 Junqi Jiang, Jianglin Lan, Francesco Leofante, Antonio Rago, and Francesca Toni. Provably robust and plausible counterfactual explanations for neural networks via robust optimisation. In Berrin Yanıkoğlu and Wray Buntine, editors, *Proceedings of the 15th Asian Conference on Machine Learning*, volume 222 of *Proceedings of Machine Learning Research*, pages 582–597. PMLR, 11–14 November 2024. URL: https://proceedings.mlr.press/v222/jiang24a.html.

31 Junqi Jiang, Francesco Leofante, Antonio Rago, and Francesca Toni. Robust counterfactual explanations in machine learning: A survey. *CoRR*, abs/2402.01928, 2024. arXiv:2402.01928.

32 Ulrich Junker. QUICKXPLAIN: preferred explanations and relaxations for over-constrained problems. In Deborah L. McGuinness and George Ferguson, editors, *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*, pages 167–172. AAAI Press / The MIT Press, 2004. URL: http://www.aaai.org/Library/AAAI/2004/aaai04-027.php.

33 Guy Katz, Clark W. Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In Rupak Majumdar and Viktor Kuncak, editors, *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I*, volume 10426 of *Lecture Notes in Computer Science*, pages 97–117. Springer, 2017. doi:10.1007/978-3-319-63387-9_5.

34 A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009.

35 Emanuele La Malfa, Rhiannon Michelmore, Agnieszka M. Zbrzezny, Nicola Paoletti, and Marta Kwiatkowska. On guaranteed optimal robust explanations for NLP models. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 2658–2665. ijcai.org, 2021. doi:10.24963/IJCAI.2021/366.

36 Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, 2, 2010.

37 Francesco Leofante and Nico Potyka. Promoting counterfactual robustness through diversity. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'24/IAAI'24/EAAI'24. AAAI Press, 2024. doi:10.1609/aaai.v38i19.30127.

38 Jianlin Li, Jiangchao Liu, Pengfei Yang, Liqian Chen, Xiaowei Huang, and Lijun Zhang. Analyzing deep neural networks with symbolic propagation: Towards higher precision and faster verification. In Bor-Yuh Evan Chang, editor, *Static Analysis - 26th International Symposium, SAS 2019, Porto, Portugal, October 8-11, 2019, Proceedings*, volume 11822 of *Lecture Notes in Computer Science*, pages 296–319. Springer, 2019. doi:10.1007/978-3-030-32304-2_15.

39 Alessio Lomuscio and Lalit Maganti. An approach to reachability analysis for feed-forward relu neural networks. *arXiv preprint*, 2017. arXiv:1706.07351.

40 Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 4765–4774, 2017. URL: https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html.

41 Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Neural Information Processing Systems*, 2017.

42 Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

43 João Marques-Silva. Logic-based explainability in machine learning. In Leopoldo E. Bertossi and Guohui Xiao, editors, *Reasoning Web. Causality, Explanations and Declarative Knowledge - 18th International Summer School 2022, Berlin, Germany, September 27-30, 2022, Tutorial*

Lectures, volume 13759 of *Lecture Notes in Computer Science*, pages 24–104. Springer, 2022. doi:10.1007/978-3-031-31414-8_2.

44 João Marques-Silva. Logic-based explainability: Past, present and future. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation. Software Engineering Methodologies - 12th International Symposium, ISoLA 2024, Crete, Greece, October 27-31, 2024, Proceedings, Part IV*, volume 15222 of *Lecture Notes in Computer Science*, pages 181–204. Springer, 2024. doi:10.1007/978-3-031-75387-9_12.

45 Joao Marques-Silva, Thomas Gerspacher, Martin C Cooper, Alexey Ignatiev, and Nina Narodytska. Explanations for monotonic classifiers. In Marina Meila and Tong Zhang, editors, *International Conference on Machine Learning*. PMLR, 2021.

46 João Marques-Silva and Alexey Ignatiev. Delivering trustworthy AI through formal XAI. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 12342–12350. AAAI Press, 2022. doi:10.1609/AAAI.V36I11.21499.

47 Mark Niklas Müller, Gleb Makarchuk, Gagandeep Singh, Markus Püschel, and Martin T. Vechev. PRIMA: general and precise neural network certification via scalable convex hull approximations. *Proc. ACM Program. Lang.*, 6(POPL):1–33, 2022. doi:10.1145/3498704.

48 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems*, 2019.

49 Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.

50 Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016*, pages 1135–1144. ACM, 2016. doi:10.1145/2939672.2939778.

51 Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic explanations. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1527–1535. AAAI Press, 2018. doi:10.1609/AAAI.V32I1.11491.

52 Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *IEEE International Conference on Computer Vision*, 2017.

53 Andy Shih, Arthur Choi, and Adnan Darwiche. A symbolic approach to explaining bayesian network classifiers. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5103–5111. ijcai.org, 2018. doi:10.24963/IJCAI.2018/708.

54 Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint*, 2013. arXiv:1312.6034.

55 Gagandeep Singh, Rupanshu Ganvir, Markus Püschel, and Martin T. Vechev. Beyond the single neuron convex barrier for neural network certification. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC,*

Canada, pages 15072–15083, 2019. URL: https://proceedings.neurips.cc/paper/2019/hash/0a9fdbb17feb6ccb7ec405cfb85222c4-Abstract.html.

56 Gagandeep Singh, Timon Gehr, Matthew Mirman, Markus Püschel, and Martin T. Vechev. Fast and effective robustness certification. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 10825–10836, 2018. URL: https://proceedings.neurips.cc/paper/2018/hash/f2f446980d8e971ef3da97af089481c3-Abstract.html.

57 Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin T. Vechev. An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.*, 3(POPL):41:1–41:30, 2019. doi:10.1145/3290354.

58 Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012. doi:10.1016/J.NEUNET.2012.02.016.

59 Vincent Tjeng, Kai Yuanqing Xiao, and Russ Tedrake. Evaluating robustness of neural networks with mixed integer programming. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL: https://openreview.net/forum?id=HyGIdiRqtm.

60 Shubham Ugare, Debangshu Banerjee, Sasa Misailovic, and Gagandeep Singh. Incremental verification of neural networks. *Proc. ACM Program. Lang.*, 7(PLDI):1920–1945, 2023. doi:10.1145/3591299.

61 Caterina Urban and Antoine Miné. A review of formal methods applied to machine learning. *CoRR*, abs/2104.02466, 2021. arXiv:2104.02466.

62 Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. Efficient formal safety analysis of neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6369–6379, 2018. URL: https://proceedings.neurips.cc/paper/2018/hash/2ecd2bd94734e5dd392d8678bc64cdab-Abstract.html.

63 Shiqi Wang, Huan Zhang, Kaidi Xu, Xue Lin, Suman Jana, Cho-Jui Hsieh, and J. Zico Kolter. Beta-crown: Efficient bound propagation with per-neuron split constraints for neural network robustness verification. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 29909–29921, 2021. URL: https://proceedings.neurips.cc/paper/2021/hash/fac7fead96dafceaf80c1daffeae82a4-Abstract.html.

64 Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, 2018.

65 Min Wu, Xiaofu Li, Haoze Wu, and Clark W. Barrett. Better verified explanations with applications to incorrectness and out-of-distribution detection. *CoRR*, abs/2409.03060, 2024. doi:10.48550/arXiv.2409.03060.

66 Min Wu, Haoze Wu, and Clark W. Barrett. Verix: Towards verified explainability of deep neural networks. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL: http://papers.nips.cc/paper_files/paper/2023/hash/46907c2ff9fafd618095161d76461842-Abstract-Conference.html.

67 Kaidi Xu, Huan Zhang, Shiqi Wang, Yihan Wang, Suman Jana, Xue Lin, and Cho-Jui Hsieh. Fast and complete: Enabling complete neural network verification with rapid and massively parallel incomplete verifiers. In *9th International Conference on Learning Representations,*

ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021. URL: https://openreview.net/forum?id=nVZtXBI6LNn.

68 Huan Zhang, Shiqi Wang, Kaidi Xu, Linyi Li, Bo Li, Suman Jana, Cho-Jui Hsieh, and J Zico Kolter. General cutting planes for bound-propagation-based neural network verification. In *Neural Information Processing Systems*, volume 35, 2022.

69 Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 4944–4953, 2018. URL: https://proceedings.neurips.cc/paper/2018/hash/d04863f100d59b3eb688a11f95b0ae60-Abstract.html.

70 Duo Zhou, Christopher Brix, Grani A. Hanasusanto, and Huan Zhang. Scalable neural network verification with branch-and-bound inferred cutting planes. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL: http://papers.nips.cc/paper_files/paper/2024/hash/33d93e4dc57453e7667b20f62e7c0681-Abstract-Conference.html.

71 Duo Zhou, Christopher Brix, Grani A Hanasusanto, and Huan Zhang. Scalable neural network verification with branch-and-bound inferred cutting planes. In *Neural Information Processing Systems*, pages 29324–29353, 2024.