# Formal Verification of Machine Learning 

MPRI 2-6: Abstract Interpretation, Application to Verification and Static Analysis


## Machine Learning Revolution

Computer software able to efficiently and autonomously perform tasks that are difficult or even impossible to design using explicit programming


Examples: object recognition, image classification, speech recognition, etc.

## ML in Safety-Critical Applications

Enables new functions that could not be envisioned before


Self-Driving Cars


Image-Based Taxiing, Takeoff, Landing

## ML in Safety-Critical Applications

Approximates complex systems and automates decision-making


Diagnosis and Drug Discovery

Deep Neural Network Compression for Aircraft Collision Avoidance Systems

Kyle D. Julian ${ }^{1}$ and Mykel J. Kochenderfer ${ }^{2}$ and Michael P. Owen ${ }^{3}$ bstract-One approach to designing decision making logic for Abstract-One app avoidance system frames the problem as a an arkov decision process and optimizes avoidance strategy can be programming. The resulting colle. This methodology has been used X propresented as a numeric Airborne Collision Avoidance System X in the development of the Airborne Colisision Avs for manned and (ACAS $\mathbf{X}$ ) family of cons the high dimensionality of the saty a deep unmanned aircraft, but the high dimens storage efficiency, a deep
. echnique to reduce the size
floating point storage. A simple technique to floa the score table is to downsample the table acter den quality, of the score ing. To minimize the degradation in program ere removed in areas where the variaiong reduces the size states are rem are smooth. The downsampling duced by dynamic in the table are smootor of 180 from that produced by dysampled of the table by ar the rest of this paper, the downsaseline, programming. For thal table is referred to as the baseline ACAS Xu ho

## ML in Safety-Critical Applications

## STAT+

IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show

By Casey Ross ${ }^{3}$ @caseymross ${ }^{4}$ and Ike Swetlitz
July 25, 2018

## A self-driving Uber ran a red light last December, contrary to company claims

Internal documents reveal that the car was at fault
By Andrew Liptak \| @AndrewLiptak \| Feb 25, 2017, 11:08am EST

## Feds Say Self-Driving Uber SUV Did Not Recognize Jaywalking Pedestrian In Fatal Crash

Richard Gonzales November 7, 201910:57 PM ET


## Machine Learning Pipeline

## DATA SCIENCE SOFTWARE


model deployment


## Machine Learning Pipeline

## Model Training is Highly Non-Deterministic



## Machine Learning Pipeline

## Models Only Give Probabilistic Guarantees



## Formal Methods

## Mathematical Guarantees of Safety



## Deductive Verification

- extremely expressive
- relies on the user to guide the proof



## Model Checking

- analysis of a model of the software
with respect to the model


Static Analysis

- analysis of the software
at some level of abstraction
- fully
and
by construction
- generally not complete


# Formal Methods for Trained Models 



## Neural Networks

## Feed-Forward Fully-Connected Neural Networks

 with ReLU Activation FunctionsRectified Linear Unit (ReLU)


## Feed-Forward Fully-Connected ReLU Networks as Programs



$$
\begin{aligned}
& x 00=\operatorname{input}() \\
& x 01=\text { input }()
\end{aligned}
$$

```
x10 = -0.31 * x00 + 0.99 * x01 + (-0.63)
x11 = -1.25 * x00 + (-0.64) * x01 + 1.88
x10 = 0 if x10 < 0 else x10
x11 = 0 if }\times11<0\mathrm{ else }\times1
x20 = 0.40 * x10 + 1.21 * x11 + 0.00
x21 = 0.64* x10 + 0.69 * x11 + (-0.39)
```

x20 $=0$ if x20<0 else x20
$\times 21=0$ if $\times 21<0$ else $\times 21$
x30 $=\mathbf{0 . 2 6}$ * $\times 20+\mathbf{0 . 3 3}$ * x21 + $\mathbf{0 . 4 5}$
X31 $=\mathbf{1 . 4 2}$ * $\times 20+\mathbf{0 . 4 0}$ * $\times 21+(-\mathbf{0 . 4 5})$
return
if $x 31<30$ else

## Maximal Trace Semantics




## Collecting Semantics



## Stability

Goal G3 in [Kurd03]

## Safety

Goal G4 in [Kurd03]

## Fairness




## Stability

Goal G3 in [Kurd03]

## Safety

Goal G4 in [Kurd03]

## Fairness


$\equiv$ Google Translate
::
Sign in
$\bar{X}_{A}$ Text Documents

ENGLISH

A nurse
A doctor
4)

16/5000 ■ - -


## Local Stability

The classification is unaffected by small input perturbations


## Local Stability

## Distance-Based Perturbations

$P_{\delta, \epsilon}(\mathbf{x}) \stackrel{\text { def }}{=}\left\{\mathbf{x}^{\prime} \in \mathscr{R}^{\left|L_{0}\right|} \mid \delta\left(\mathbf{x}, \mathbf{x}^{\prime}\right) \leq \epsilon\right\}$
Example ( $L_{\infty}$ distance): $P_{\infty, \epsilon}(\mathbf{x}) \stackrel{\text { def }}{=}\left\{\mathbf{x}^{\prime} \in \mathscr{R}^{\left|L_{0}\right|}\left|\max _{i}\right| \mathbf{x}_{i}-\mathbf{x}_{i}^{\prime} \mid \leq \epsilon\right\}$

$$
\mathscr{R}_{\mathbf{x}}^{\delta, \epsilon} \stackrel{\text { def }}{=}\left\{\llbracket M \rrbracket \in \mathscr{P}\left(\Sigma^{*}\right) \mid \operatorname{STABLE}_{\mathbf{x}}^{\delta, \epsilon}(\llbracket M \rrbracket)\right\}
$$

$\mathscr{R}_{\mathbf{x}}^{\delta, \epsilon}$ is the set of all neural networks M (or, rather, their semantics $\left.\llbracket M \rrbracket\right)$ that are stable in the neighborhood $P_{\delta, \epsilon}(\mathbf{x})$ of a given input $\mathbf{x}$
$\operatorname{STABLE}_{\mathbf{x}}^{\delta, \epsilon}(\llbracket M \rrbracket) \stackrel{\text { def }}{=} \forall t \in \llbracket M \rrbracket:\left(\exists t^{\prime} \in \llbracket M \rrbracket: \forall 0 \leq i \leq\left|L_{0}\right|: t_{0}^{\prime}\left(x_{0, i}\right)=\mathbf{x}_{i}\right)$

$$
\begin{aligned}
& \wedge\left(\exists \mathbf{x}^{\prime} \in P_{\delta, \epsilon}(\mathbf{x}): \forall 0 \leq i \leq\left|L_{0}\right|: t_{0}\left(x_{0, i}\right)=\mathbf{x}_{i}^{\prime}\right) \\
& \Rightarrow \max _{j} t_{\omega}\left(x_{N, j}\right)=\max _{j} t_{\omega}^{\prime}\left(x_{N, j}\right)
\end{aligned}
$$

## Theorem

## Corollary

$$
M \vDash \mathscr{R}_{\mathbf{x}}^{\delta, \epsilon} \Leftrightarrow \llbracket M \rrbracket \subseteq \bigcup \mathscr{R}_{\mathbf{x}}^{\delta, \epsilon}
$$



## Forward Analysis



## Example


$P(\langle 0.5,0.75\rangle) \stackrel{\text { def }}{=}\left\{\mathbf{x} \in \mathscr{R} \times \mathscr{R} \mid 0 \leq \mathbf{x}_{0} \leq 1 \wedge 0 \leq \mathbf{x}_{1} \leq 1\right\}$
$x_{i, j} \mapsto[a, b]$
$a, b \in \mathscr{R}$


## Interval Domain

 combination of the inputs with Symbolic Constant Propagation and the previous ReLUs$x_{i, j} \mapsto \begin{cases}\sum_{k=0}^{i-1} \mathbf{c}_{k} \cdot \mathbf{x}_{k}+\mathbf{c} & \mathbf{c}_{k}, \mathbf{c} \in \mathscr{R}^{\left|\mathbf{X}_{k}\right|} \\ {[a, b]} & a, b \in \mathscr{R}\end{cases}$


$$
\begin{aligned}
& x_{i, j} \mapsto\left\{\begin{array}{l}
\mathbf{E}_{\mathbf{i}, \mathbf{j}} \\
{[\mathrm{a}, \mathrm{~b}]}
\end{array} \quad 0 \leq a\right. \\
& \begin{aligned}
x_{i, j} \mapsto\left\{\begin{array} { l } 
{ \mathbf { E } _ { \mathbf { i } , \mathbf { j } } } \\
{ [ \mathbf { a } , \mathbf { b } ] } \\
{ }
\end{array} x _ { i , j } \mapsto \left\{\begin{array}{l}
\mathbf{x}_{\mathbf{i}, \mathbf{j}} \\
{[0, \mathrm{~b}]}
\end{array}\right.\right. & a<0 \wedge 0<b
\end{aligned}
\end{aligned}
$$

J. Li et al. - Analyzing Deep Neural Networks with Symbolic Propagation (SAS 2019)

## Interval Domain

## with Symbolic Constant Propagation [ui9]



## Interval Domain

 with Symbolic Constant Propagation

## DeepPoly ${ }_{\text {Binnve }}$

maintain symbolic lower- and upper-bounds for each neuron + convex ReLU approximations
$x_{i+1, j} \mapsto \begin{cases}{\left[\sum_{k} c_{i, k} \cdot x_{i, k}+c, \sum_{k} d_{i, k} \cdot x_{i, k}+d\right]} & c_{i, k} c, d_{i, k}, d \in \mathscr{R} \\ {[a, b]} & a, b \in \mathscr{R}\end{cases}$



G. Singh, T. Gehr, M. Püschel, and M. Vechev - An Abstract Domain for Certifying Neural Networks (POPL 2019)

## DeepPoly ${ }_{\text {Binnve }}$



## DeepPoly



## DeepPoly



## DeepPoly ${ }_{\text {Binnve }}$



$$
x_{40} \mapsto\left\{\left[0.5 \cdot x_{30}-2 \cdot x_{31}+1,0.5 \cdot x_{30}-2 \cdot x_{31}+1\right]\right.
$$

## DeOBPOM [Singh19]

$$
\begin{aligned}
x_{00} & \mapsto \begin{cases}{\left[x_{00}, x_{00}\right]} \\
{[\mathbf{0}, \mathbf{1}]}\end{cases}
\end{aligned} x_{01} \mapsto\left\{\begin{array}{l}
{\left[x_{01}, x_{01}\right]} \\
{[\mathbf{0}, \mathbf{1}]}
\end{array}\right]
$$

## DeepPoly ${ }_{\text {Binnve }}$

$$
x_{40} \mapsto\left\{\begin{array}{l}
{\left[0.5 \cdot x_{30}-2 \cdot x_{31}+1,0.5 \cdot x_{30}-2 \cdot x_{31}+1\right]} \\
{[\underline{\mathbf{2}}, \mathbf{5}]}
\end{array}\right.
$$



## DePOBOM [Singh19]

$$
\begin{aligned}
x_{00} & \mapsto \begin{cases}{\left[x_{00}, x_{00}\right]} \\
{[\mathbf{0}, \mathbf{1}]}\end{cases}
\end{aligned} x_{01} \mapsto\left\{\begin{array}{l}
{\left[x_{01}, x_{01}\right]} \\
{[\mathbf{0}, \mathbf{1}]}
\end{array}\right\}
$$

## DeepPoly

$$
x_{40} \mapsto\left\{\begin{array}{l}
{\left[0.5 \cdot x_{30}-2 \cdot x_{31}+1,0.5 \cdot x_{30}-2 \cdot x_{31}+1\right]} \\
{[\underline{\mathbf{2}}, \mathbf{5}]}
\end{array}\right.
$$



$$
\begin{aligned}
& x_{00} \mapsto\left\{\begin{array}{l}
{\left[x_{00}, x_{00}\right]} \\
{[\mathbf{0}, \mathbf{1}]}
\end{array} O_{\sigma}\right. \\
& x_{01} \mapsto\left\{\begin{array}{l}
{\left[x_{01}, x_{01}\right]} \\
{[\mathbf{0}, \mathbf{1}]}
\end{array}\right.
\end{aligned}
$$



## Other Static Analysis Methods

- T. Gehr, M. Mirman, D. Drachsler-Cohen, P. Tsankov, S. Chaudhuri, and M. Vechev. Al2: Safety and Robustness Certification of Neural Networks with Abstract Interpretation. In S\&P, 2018. the first use of abstract interpretation for verifying neural networks
- G. Singh, T. Gehr, M. Mirman, M. Püschel, and M. Vechev. Fast and Effective Robustness Certification. In NeurIPS, 2018. a custom zonotope domain for certifying neural networks
- G. Singh, R. Ganvir, M. Püschel, and M. Vechev. Beyond the Single Neuron Convex Barrier for Neural Network Certification. In NeurIPS, 2019. a framework to jointly approximate k ReLU activations
- M. N. Müller, G. Makarchuk, G. Singh, M. Püschel, and M. Vechev. PRIMA: General and Precise Neural Network Certification via Scalable Convex Hull Approximations. In POPL, 2022.
a multi-neuron abstraction via a convex-hull approximation algorithm


## Stability

Goal G3 in [Kurd03]

## Safety

Goal G4 in [Kurd03]


## Fairness

## Airborne Collision Avoidance System for Unmanned Aircraft

implemented using 45 feed-forward fully-connected ReLU networks


## 5 input sensor measurements

- $\rho$ : distance from ownship to intruder
- $\theta$ : angle to intruder relative to ownship heading direction
- $\psi$. heading angle to intruder relative to ownship heading direction
- $v_{\text {own }}$ : speed of ownship
- $v_{\text {int }}$ : speed of intruder


## 5 output horizontal advisories



- Strong Left
- Weak Left
- Clear of Conflict
- Weak Right
- Strong Right


## ACAS Xu Properties

Example: "if intruder is near and approaching from the left, go Strong Right"


## Safety

## Input-Output Properties

I: input specification
O: output specification

$$
\mathcal{S}_{\mathbf{O}}^{\mathbf{I}} \stackrel{\text { def }}{=}\left\{\llbracket M \rrbracket \in \mathscr{P}\left(\Sigma^{*}\right) \mid \operatorname{SAFE}_{\mathbf{O}}^{\mathbf{I}}(\llbracket M \rrbracket)\right\}
$$

$\mathcal{S}_{\mathbf{O}}^{\mathbf{I}}$ is the set of all neural networks M (or, rather, their semantics $\left.\llbracket M \rrbracket\right)$ that satisfy the input and output specification $\mathbf{I}$ and $\mathbf{O}$ $\operatorname{SAFE}_{\mathbf{O}}^{\mathbf{I}}(\llbracket M \rrbracket) \stackrel{\text { def }}{=} \forall t \in \llbracket M \rrbracket: t_{0} \vDash \mathbf{I} \Rightarrow t_{\omega} \vDash \mathbf{O}$

## Theorem

$M \leftarrow \delta_{0}^{1} \Leftrightarrow\{\|M\|\} \subseteq \delta_{0}^{1}$

Corollary

$$
M \vDash \delta_{\mathbf{O}}^{\mathbf{I}} \Leftrightarrow \llbracket M \rrbracket \subseteq \bigcup \delta_{\mathbf{O}}^{\mathbf{I}}
$$

Formal Methods
Mathematical
Mathematical Guarantees


Deductive Verification
relies on the usersive

Model Checking
analysis of a mocking
sound and complete the software
with respect to the te

Static Analysis
analysis Analy
at some level software
fully automatic and abstraction generally not complete sound by
ormal Verification of Machine Learning
Calerina urban

## Model Checking Methods

## Safety

## Example



## SMT-Based Methods

## Verification Reduced to Constraint Satisfiability

$$
\begin{array}{ll}
\mathbf{l}_{\mathrm{j}} \leq \mathbf{x}_{0, \mathrm{j}} \leq \mathbf{u}_{\mathbf{j}} & j \in\left\{0, \ldots,\left|\mathbf{X}_{0}\right|\right\} \\
\hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{x}_{i}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} & i \in\{0, \ldots, n-1\} \\
x_{i, j}=\max \left\{0, \hat{x}_{i, j}\right\} & i \in\{1, \ldots, n-1\}, \\
j \in\left\{0, \ldots,\left|\mathbf{X}_{i}\right|\right\}
\end{array}
$$

$\mathbf{x}_{\mathrm{N}} \leq \mathbf{0}$
input specification

(negation of) output specification
satisfiable $\rightarrow$ X counterexample otherwise $\rightarrow \boldsymbol{\downarrow}$ safe

## Planet

## use approximations to reduce the solution search

$$
\mathbf{0} \leq \mathbf{x}_{\mathrm{i}, \mathrm{j}}
$$

$x_{i, j}=\max \left\{0, \hat{x}_{i, j}\right\}$

$$
\begin{aligned}
& \hat{\mathbf{x}}_{\mathbf{i}, \mathbf{j}} \leq \mathbf{x}_{\mathbf{i}, \mathbf{j}} \\
& \mathbf{x}_{\mathbf{i}, \mathbf{j}} \leq \frac{\mathbf{b}_{\mathbf{i}, \mathbf{j}}}{\mathbf{b}_{\mathbf{i}, \mathbf{j}}-\mathbf{a}_{\mathbf{i}, \mathbf{j}}} \cdot\left(\hat{\mathbf{x}}_{\mathbf{i}, \mathbf{j}}-\mathbf{a}_{\mathbf{i}, \mathbf{j}}\right)
\end{aligned}
$$



[^0]
## Reluplex

## based on the simplex algorithm extended to support ReLUs

| Variable | Value |
| :---: | :---: |
| $\mathbf{x}_{\mathbf{0 0}}$ | $v_{00}$ |
| $\cdots$ | $\cdots$ |
| $\hat{\mathbf{x}}_{\mathbf{i}, \mathrm{j}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\mathbf{X}_{\mathrm{i} \mathrm{j}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\cdots$ | $\cdots$ |
| $\mathbf{X}_{\mathbf{N}}$ | $v_{N}$ |


| Variable | Value |
| :---: | :---: |
| $\mathbf{x}_{\mathbf{0 0}}$ | $v_{00}$ |
| $\cdots$ | $\cdots$ |
| $\hat{\mathbf{x}}_{\mathbf{i}, \mathrm{j}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\mathbf{x}_{\mathrm{i} \mathrm{j}}$ | 0 |
| $\cdots$ | $\cdots$ |
| $\mathbf{X}_{\mathbf{N}}$ | $v_{N}$ |

G. Katz et al. - Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks (CAV 2017)

## Reluplex

based on tr extended
E

| Variable | Value |
| :---: | :---: |
| $\mathbf{x}_{\mathbf{0 0}}$ | $v_{00}$ |
| $\cdots$ | $\cdots$ |
| $\hat{\mathbf{x}}_{\mathbf{i}, \mathrm{j}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\mathbf{X}_{\mathbf{i j}}$ | $v_{i j}$ |
| $\cdots$ | $\cdots$ |
| $\mathbf{X}_{\mathbf{N}}$ | $v_{N}$ |

## Variable

| Variable | Value |
| :---: | :---: |
| $\mathbf{x}_{\mathbf{0 0}}$ | $v_{00}$ |
| $\cdots$ | $\cdots$ |
| $\hat{\mathbf{x}}_{\mathbf{i} \mathbf{j}}$ | $\hat{v}_{i j}$ |
| $\mathbf{X}_{\mathbf{i j}}$ | $v_{i j}$ |
| $\cdots$ | $\cdots$ |
| $\mathbf{x}_{\mathbf{N}}$ | $v_{N}$ |

G. Katz et al. - The

Marabou Framework for Verification and Analysis of Deep Neural Networks (CAV 2019)

| $\mathbf{x}_{\mathbf{0 0}}$ | $v_{00}$ |
| :---: | :---: |
| $\cdots$ | $\cdots$ |
| $\hat{\mathbf{x}}_{\mathrm{i}, \mathrm{j}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\mathbf{X}_{\mathrm{ij}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\cdots$ | $\cdots$ |
| $\mathbf{X}_{\mathbf{N}}$ | $v_{N}$ |


| Variable | Value |
| :---: | :---: |
| $\mathbf{X}_{\mathbf{0 0}}$ | $v_{00}$ |
| $\cdots$ | $\cdots$ |
| $\hat{\mathbf{x}}_{\mathbf{i}, \mathrm{j}}$ | $\hat{v}_{i j}^{\prime}$ |
| $\mathbf{X}_{\mathrm{ij}}$ | 0 |
| $\cdots$ | $\cdots$ |
| $\mathbf{X}_{\mathbf{N}}$ | $v_{N}$ |

G. Katz et al. - Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks (CAV 2017)

## Other SMT-Based Methods

- L. Pulina and A. Tacchella. An Abstraction-Refinement Approach to Verification of Artificial Neural Networks. In CAV, 2010. the first formal verification method for neurall networks
- O. Bastani, Y. Ioannou, L. Lampropoulos, D. Vytiniotis, A. Nori, and A. Criminisi. Measuring Neural Net Robustness with Constraints. In NeurIPS, 2016. an approach for finding the nearest adversarial example according to the Lo distance
- X. Huang, M. Kwiatkowska, S. Wang, and M. Wu. Safety Verification of Deep Neural Networks. In CAV, 2017. an approach for proving local robustness to adversarial perturbations
- N. Narodytska, S. Kasiviswanathan, L. Ryzhyk, M. Sagiv, and T. Walsh. Verifying Properties of Binarized Deep Neural Networks. In AAAI, 2018. C. H. Cheng, G. Nührenberg, C. H. Huang, and H. Ruess. Verification of Binarized Neural Networks via Inter-Neuron Factoring. In VSTTE, 2018. approaches focusing on binarized neural networks


## MILP-Based Methods

## Verification Reduced to Mixed Integer Linear Program

$$
\begin{array}{ll}
\mathrm{I}_{\mathrm{j}} \leq \mathrm{x}_{0, \mathrm{j}} \leq \mathbf{u}_{\mathrm{j}} & j \in\left\{0, \ldots,\left|\mathbf{X}_{0}\right|\right\} \\
\hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{x}_{\mathrm{i}}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} & i \in\{0, \ldots, n-1\} \\
x_{i, j}=\delta_{\mathrm{i}, \mathrm{j}} \cdot \hat{x}_{i, j} & \delta_{\mathrm{i}, \mathrm{j}} \in\{\mathbf{0 , 1}\} \\
\delta_{i, \mathrm{j}}=1 \Rightarrow \hat{x}_{i, j} \geq 0 & i \in\{1, \ldots, n-1\} \\
\delta_{\mathrm{i}, \mathrm{j}}=0 \Rightarrow \hat{x}_{i, j}<0 & j \in\left\{0, \ldots,\left|\mathbf{X}_{i}\right|\right\}
\end{array}
$$

$\min \mathbf{x}_{\mathrm{N}}$

objective function


## MILP-Based Methods

## Bounded Encoding with Symmetric Bounds

$$
\begin{array}{ll}
\hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{X}_{i}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} & i \in\{0, \ldots, n-1\} \\
0 \leq x_{i, j} \leq \mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot \delta_{i, j} & \delta_{\mathbf{i}, \mathbf{j}} \in\{\mathbf{0}, \mathbf{1}\} \\
\hat{x}_{i, j} \leq x_{i, j} \leq \hat{x}_{i, j}-\mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot\left(1-\delta_{i, j}\right) & i \in\{1, \ldots, n-1\} \\
\mathbf{M}_{\mathbf{i}, \mathrm{j}}=\max \left\{-\mathbf{l}_{\mathbf{i}}, \mathbf{u}_{\mathbf{i}}\right\} & j \in\left\{0, \ldots,\left|\mathbf{X}_{i}\right|\right\}
\end{array}
$$



## Sherlock

## Output Range Analysis

## use local search to speed up the MILP solver

$$
\begin{aligned}
& \mathbb{1}_{\mathrm{j}} \leq \mathbf{x}_{\mathbf{0 , j}} \leq \mathbf{u}_{\mathbf{j}} \\
& \hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{x}_{i}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} \\
& 0 \leq x_{i, j} \leq \mathbf{M}_{\mathbf{i}, \mathrm{j}} \cdot \delta_{i, j} \\
& \hat{x}_{i, j} \leq x_{i, j} \leq \hat{x}_{i, j}-\mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot\left(1-\delta_{i, j}\right) \\
& \mathbf{M}_{\mathbf{i}, \mathbf{j}}=\max \left\{-\mathbf{l}_{\mathbf{i}}, \mathbf{u}_{\mathbf{i}}\right\} \\
& \mathbf{x}_{\mathbf{N}}<\mathbf{L}
\end{aligned}
$$

[^1]
## Sherlock

## Output Range Analysis

## use local search to speed up the MILP solver

$$
\begin{aligned}
& \mathbf{1}_{\mathbf{j}} \leq \mathbf{x}_{0, \mathbf{j}} \leq \mathbf{u}_{\mathbf{j}} \\
& \hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{X}_{i}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} \\
& 0 \leq x_{i, j} \leq \mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot \delta_{i, j} \\
& \hat{x}_{i, j} \leq x_{i, j} \leq \hat{x}_{i, j}-\mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot\left(1-\delta_{i, j}\right) \\
& \mathbf{M}_{\mathbf{i}, \mathbf{j}}=\max \left\{-\mathbf{l}_{\mathbf{i}}, \mathbf{u}_{\mathbf{i}}\right\} \\
& \mathbf{x}_{\mathbf{N}}<\hat{\mathbf{L}}
\end{aligned}
$$

[^2]
## Sherlock

## Output Range Analysis

## use local search to speed up the MILP solver

$$
\begin{aligned}
& \boldsymbol{1}_{\mathbf{j}} \leq \mathbf{x}_{\mathbf{0}, \mathbf{j}} \leq \mathbf{u}_{\mathbf{j}} \\
& \hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{X}_{i}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} \\
& 0 \leq x_{i, j} \leq \mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot \delta_{i, j} \\
& \hat{x}_{i, j} \leq x_{i, j} \leq \hat{x}_{i, j}-\mathbf{M}_{\mathbf{i}, \mathbf{j}} \cdot\left(1-\delta_{i, j}\right) \\
& \mathbf{M}_{\mathbf{i}, \mathbf{j}}=\max \left\{-\mathbf{l}_{\mathbf{i}}, \mathbf{u}_{\mathbf{i}}\right\} \\
& \mathbf{x}_{\mathbf{N}}<\hat{\mathbf{L}}
\end{aligned}
$$


find another input $\hat{\mathbf{X}}$ such that $\hat{\mathbf{L}} \leq \mathbf{x}_{\mathbf{N}}$
S. Dutta et al. - Output Range Analysis for Deep Feedforward Neural Networks (NFM 2018)

## MILP-Based Methods

## Bounded Encoding with Asymmetric Bounds

$$
\begin{array}{ll}
\hat{x}_{i+1, j}=\sum_{k=0}^{\left|\mathbf{X}_{\mathbf{N}}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} & i \in\{0, \ldots, n-1\} \\
0 \leq x_{i, j} \leq \mathbf{u}_{i, j} \cdot \delta_{i, j} & \delta_{i, \mathbf{j}} \in\{\mathbf{0}, \mathbf{1}\} \\
\hat{x}_{i, j} \leq x_{i, j} \leq \hat{x}_{i, j}-\mathbf{l}_{\mathbf{i}, \mathrm{j}} \cdot\left(1-\delta_{i, j}\right) & i \in\{1, \ldots, n-1\} \\
& j \in\left\{0, \ldots,\left|\mathbf{X}_{i}\right|\right\}
\end{array}
$$



## MIPVerify

## Finding Nearest Adversarial Example

$$
\mathbf{x}_{\mathbf{N}} \neq \mathbf{O}
$$

V. Tjeng et al. - Evaluating Robustness of Neural Networks with Mixed Integer Programming (ICLR 2019)

$$
\begin{aligned}
& \min _{\mathbf{X}^{\prime}} \mathbf{d}\left(\mathbf{X}, \mathbf{X}^{\prime}\right) \\
& \hat{x}_{i+1, j}=\sum_{k=0}^{\left|X_{\mid}\right|} w_{j, k}^{i} \cdot x_{i, k}+b_{i, j} \quad i \in\{0, \ldots, n-1\} \\
& 0 \leq x_{i, j} \leq \mathbf{u}_{\mathrm{i}, \mathrm{j}} \cdot \delta_{i, j} \\
& \delta_{\mathrm{i}, \mathrm{j}} \in\{\mathbf{0}, \mathbf{1}\} \\
& \hat{x}_{i, j} \leq x_{i, j} \leq \hat{x}_{i, j}-\mathbf{l}_{\mathbf{i}, \mathrm{j}} \cdot\left(1-\delta_{i, j}\right) \quad \begin{array}{l}
i \in\{1, \ldots, n-1\} \\
j \in\left\{0, \ldots,\left|\mathbf{X}_{i}\right|\right\}
\end{array}
\end{aligned}
$$

## Other MILP-Based Methods

- R. Bunel, I. Turkaslan, P. H. S. Torr, P. Kohli, and M. P. Kumar. A Unified View of Piecewise Linear Neural Network Verification. In NeurlPS, 2018. a unifying verification framework for piecewise-linear ReLU neural networks
- C.-H. Cheng, G. Nührenberg, and H. Ruess. Maximum Resilience of Artificial Neural Networks. In ATVA, 2017. an approach for finding a lower bound on robustness to adversarial perturbations
- M. Fischetti and J. Jo. Deep Neural Networks and Mixed Integer Linear Optimization. 2018. an approach for feature visualization and building adversarial examples



## Static Analysis Methods

## Forward Analysis

(2) check output for inclusion in output specification $\mathbf{O}$ : included $\rightarrow$ safe otherwise $\rightarrow$ alarm
(1) proceed forwards from an abstraction of the input specification I

## Example



## DeepPoly Domain ${ }_{\text {Einnve }}$

$\begin{aligned} & \operatorname{ReLU}\left(\begin{array}{l}x_{10}\end{array} \mapsto\left\{\begin{array}{l}{\left[x_{10}, \frac{2}{3} \cdot x_{10}+\frac{2}{3}\right]} \\ {[-1,2]}\end{array}\right.\right. \\ & x_{10} \mapsto\left\{\begin{array}{l}{\left[x_{00}+x_{01}, x_{00}+x_{01}\right]} \\ {[-1,2]}\end{array}\right.\end{aligned}$

$$
\begin{aligned}
& 0 \leq \rho \leq 1
\end{aligned}
$$




## DeepPoly Domain ${ }_{\text {Einnve }}$



## DeepPoly Domain ${ }_{\text {sirnvel }}$



## Interval Domain

## with Symbolic Constant Propagation



## Interval Domain

## with Symbolic Constant Propagation



## Interval Domain

## with Symbolic Constant Propagation [Lit9]



## Product Domain ${ }_{\text {weazacaren }}$

$$
\begin{aligned}
& x_{00} \mapsto\left\{\begin{array}{l}
x_{00} \\
{\left[x_{00}, x_{00}\right]} \\
{[\mathbf{0}, \mathbf{1}]}
\end{array}\right. \\
& x_{01} \mapsto\left\{\begin{array}{l}
x_{01} \\
{\left[x_{01}, x_{01}\right]} \\
{[-\mathbf{1}, \mathbf{1}]}
\end{array}\right. \\
& -1 \leq \theta \leq 1 \\
& \times 00
\end{aligned}
$$

## Product Domain ${ }_{\text {Mazzucatar1] }}$

$$
x_{20} \mapsto \begin{cases}x_{10}+x_{11} & \rightarrow[\mathbf{0}, \mathbf{4}] \\ {\left[x_{10}+x_{11}, x_{10}+x_{11}\right]} & \rightarrow\left[\mathbf{0}, \frac{8}{3}\right] \\ {\left[\mathbf{0}, \frac{8}{3}\right]} & \end{cases}
$$



## Drocuctan Don Rin [Mazzucato21]

$$
x_{30} \mapsto \begin{cases}x_{10}+x_{11}+x_{21}+1 & \rightarrow\left[\mathbf{1}, \frac{\mathbf{2 0}}{\mathbf{3}}\right] \\ {\left[x_{20}+x_{21}+1, x_{20}+x_{21}+1\right]} & \rightarrow[\mathbf{1}, \mathbf{4} .5] \\ {[\underline{\mathbf{1}}, \mathbf{4} . \mathbf{5}]} & \end{cases}
$$

$$
\begin{aligned}
& 0 \leq \rho \leq 1 \\
x_{00} & \mapsto\left\{\begin{array}{llll}
x_{00} \\
{\left[x_{00}, x_{00}\right]} \\
{[\mathbf{0}, \mathbf{1}]}
\end{array}\right. \\
x_{01} \mapsto\left\{\begin{array}{lll}
x_{01} \\
{\left[x_{01}, x_{01}\right]} \\
{[-\mathbf{1}, \mathbf{1}]}
\end{array}\right. & \times 00
\end{aligned}
$$



## Exact Static Analysis Method

$V=\left\{v_{1}, \ldots, v_{m}\right\}$ : basis vectors in $\mathscr{R}^{n}$
$P: \mathscr{R}^{m} \rightarrow\{\perp, \top\}:$ predicate


$$
\llbracket \Theta \rrbracket=\left\{x \mid x=c+\sum_{i=1}^{m} \alpha_{i} v_{i} \text { such that } P\left(\alpha_{1}, \ldots, \alpha_{m}\right)=\mathrm{\top}\right\}
$$

- fast and cheap affine mapping operations $\rightarrow$ neural network layers
- inexpensive intersections with half-spaces $\rightarrow$ ReLU activations

[^3]
## ReluVal

## use symbolic propagation + iterative input refinement

## Asymptotically Complete Method


S. Wang et al. - Formal Security Analysis of Neural Networks Using Symbolic Intervals (USENIX Security 2018)

## Neurify

 convex ReLU approximation +
## Asymptotically Complete Method

$x_{i, j} \mapsto \begin{cases}{\left[\sum_{k} c_{0, k} \cdot x_{0, k}+c, \sum_{k} d_{0, k} \cdot x_{0, k}+d\right]} & c_{0, k}, c, d_{0, k}, d \in \mathscr{R} \\ {[a, b]} & a, b \in \mathscr{R}\end{cases}$


## Further Complete Methods

- W. Ruan, X. Huang, and M. Kwiatkowska. Reachability Analysis of Deep Neural Networks with Provable Guarantees. In IJCAI, 2018. a global optimization-based approach for verifying Lipschitz continuous neural networks
- G. Singh, T. Gehr, M. Püschel, and M. Vechev. Boosting Robustness Certification of Neural Networks. In ICLR, 2019. an approach combining abstract interpretation and (mixed integer) linear programming



## Other Incomplete Methods

## Interval Neural Networks

## Abstraction-Based Method



## 

Y. Y. Elboher et al. - An Abstraction-Based
Framework for Neural Network Verification (CAV 2020)

[^4]
## Further Incomplete Methods

- W. Xiang, H.-D. Tran, and T. T. Johnson. Output Reachable Set Estimation and Verification for Multi-Layer Neural Networks. 2018. an approach combining simulation and linear programming
- K. Dvijotham, R. Stanforth, S. Gowal, T. Mann, and P. Kohli. A Dual Approach to Scalable Verification of Deep Networks. In UAI, 2018. an approach based on duality for verifying neural networks


## Further Incomplete Methods

- E. Wong and Z. Kolter. Provable Defenses Against Adversarial Examples via the Convex Outer Adversarial Polytope. In ICML, 2018.
A. Raghunathan, J. Steinhardt, and P. Liang. Certified Defenses against Adversarial Examples. In ICML, 2018.
T.-W. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, L. Daniel, D.

Boning, and I. Dhillon. Towards Fast Computation of Certified Robustness for ReLU Networks. In ICML, 2018.
H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel. Efficient Neural Network Robustness Certification with General Activation Functions. In NeurlPS, 2018.
approaches for finding a lower bound on robustness to adversarial perturbations

## Further Incomplete Methods

- A. Boopathy, T.-W. Weng, P.-Y. Chen, S. Liu, and L. Daniel. CNN-Cert: An Efficient Framework for Certifying Robustness of Convolutional Neural Networks. In AAAI, 2019. approach focusing on convolutional neural networks
- C.-Y. Ko, Z. Lyu, T.-W. Weng, L. Daniel, N. Wong, and D. Lin. POPQORN: Quantifying Robustness of Recurrent Neural Networks. In ICML, 2019. H. Zhang, M. Shinn, A. Gupta, A. Gurfinkel, N. Le, and N. Narodytska. Verification of Recurrent Neural Networks for Cognitive Tasks via Reachability Analysis. In ECAI, 2020. approaches focusing on recurrent neural networks
- D. Gopinath, H. Converse, C. S. Pasareanu, and A. Taly. Property Inference for Deep Neural Networks. In ASE, 2019. an approach for inferring safety properties of neural networks


## Complete Methods

## Advantages

sound and complete
suffer from false positives

## Disadvantages

## Disadvantages

able to scale to large models
soundness not typically guaranteed with respect to floating-point arithmetic
do not scale to large models
often limited to certain
less limited to certain model architectures

Advantages
Incomplete Methods

## Stability

Goal G3 in [Kurd03]

## Safety

Goal G4 in [Kurd03]


## ML Impacts Our Society

UIRED

In 2019, predictive alunrithm
Machine Bias
Will startto
There's software used across the country to predict future criminals. And it's biased against blacks.
by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica May 23, 2016

## D CHECKS ARE 1 A HOME

## Can AI Be a Fair Judge in C Estonia Thinks So

Estonia plans to use an artificial intelligence program to
small-claims cases, part of a push to make program to smarter.

## Translation tutorial:

## 21 fairness definitions and their politics

> Arvind Narayanan
> @random_walker

## $>$ - 0:05/55:20



Tutorial: 21 fairness definitions and their politics
19,759 views • Mar 1, 2018


SHARE $\overline{=}$ SAVE
-••

226 subscribers

Computer scientists and statisticians have devised numerous mathematical criteria to define what it means for a classifier or a model to be fair. The proliferation of these definitions represents an attempt to make technical sense of SHOW MORE

## Dependency Fairness

The classification is independent of the values of the sensitive inputs



## Dependency Fairness

$$
\mathscr{F}_{i} \stackrel{\text { def }}{=}\left\{\llbracket M \rrbracket \in \mathscr{P}\left(\Sigma^{*}\right) \mid \text { UNUSED }(\llbracket M \rrbracket)\right\}
$$

$\mathscr{F}_{i}$ is the set of all neural networks M (or, rather, their semantics $\left.\llbracket M \rrbracket\right)$ that do not use the value of the sensitive input node $x_{0, i}$ for classification
$\operatorname{UNUSED}_{i}(\llbracket M \rrbracket) \stackrel{\text { def }}{=} \forall t \in \llbracket M \rrbracket, v \in \mathscr{R}: t_{0}\left(x_{0, i}\right) \neq v \Rightarrow \exists t^{\prime} \in \llbracket M \rrbracket:$

$$
\begin{aligned}
& \left(\forall 0 \leq j \leq\left|L_{0}\right|: j \neq i \Rightarrow t_{0}\left(x_{0, j}\right)=t_{0}^{\prime}\left(x_{0, j}\right)\right) \\
& \wedge t_{0}^{\prime}\left(x_{0, i}\right)=v \\
& \wedge \max _{j} t_{\omega}\left(x_{N, j}\right)=\max _{j} t_{\omega}^{\prime}\left(x_{N, j}\right)
\end{aligned}
$$

Intuitively: any possible classification outcome is possible from any value of the sensitive input node $x_{0, i}$

Input Data (Non-)Usage

$\mathcal{N}_{J}$ is the set of all programs $P$ (or, rather, their semantics $\left.\|P \mathbb{P}\|\right)$
that do not use the value of the inpu

Intuitively: any possible program
outcome is possible from
of the input variable $x_{0, i}$

## Dependency Fairness



## Dependency Fairness

$$
\mathscr{F}_{i} \stackrel{\text { def }}{=}\left\{\llbracket M \rrbracket \in \mathscr{P}\left(\Sigma^{*}\right) \mid \operatorname{UNUSED}(\llbracket M \rrbracket)\right\}
$$

$\mathscr{F}_{i}$ is the set of all neural networks M (or, rather, their semantics $\left.\llbracket M \rrbracket\right)$ that do not use the value of the sensitive input node $x_{0, i}$ for classification
$\operatorname{UNUSED}_{i}(\llbracket M \rrbracket) \stackrel{\text { def }}{=} \forall t \in \llbracket M \rrbracket, v \in \mathscr{R}: t_{0}\left(x_{0, i}\right) \neq v \Rightarrow \exists t^{\prime} \in \llbracket M \rrbracket:$ $\left(\forall 0 \leq j \leq\left|L_{0}\right|: j \neq i \Rightarrow t_{0}\left(x_{0, j}\right)=t_{0}^{\prime}\left(x_{0, j}\right)\right)$ $\wedge t_{0}^{\prime}\left(x_{0, i}\right)=v$ $\wedge \max _{j} t_{\omega}\left(x_{N, j}\right)=\max _{j} t_{\omega}^{\prime}\left(x_{N, j}\right)$
Intuitively: any possible classification outcome is possible from any value of the sensitive input node $x_{0, i}$

## Theorem

$M \vDash \mathscr{F}_{i} \Leftrightarrow\{\|M\|\} \subseteq \mathscr{F}_{i}$

## Hierarchy of Semantics

parallel semantics


## Collecting Semantics <br> \section*{为}



$$
\begin{aligned}
& \text { Hence: Col(prog) } \stackrel{\text { der }}{=}\{[\text { orog I }]\} \\
& \text { Benefitis: } \\
& \text { uniformity of semantics and properties, } \subseteq \text { information order } \\
& \text { then a program prog and a property } \left.P \in \mathcal{P}\left(\sum^{*}\right)\right) \\
& \text { therification problem is an inclusion check: } \left.\mathcal{P}\left(\sum^{*}\right)\right)
\end{aligned}
$$

$$
\begin{aligned}
& \text { " generally, the collecting } \quad \text { Co/(prog) } \subseteq P \\
& \text { we sett }
\end{aligned}
$$

$$
\begin{aligned}
& \text { secleally, the collecting semantics canog } \\
& \text { we settle for a weak } \subseteq P
\end{aligned}
$$

$$
\begin{aligned}
& \text { We settle for a weecting semantics cannot be computed, } \\
& \text { - is sound: Col(prorg) property s\# that } \\
& \text { implies } \text { st }
\end{aligned}
$$

$$
\begin{aligned}
& \text { \# is sound: wol(proer property } S^{\sharp} \text { than } \\
& \text { implies the } \text { idesire) } \subseteq S^{\sharp}
\end{aligned}
$$

$$
\text { Implies the desired property: } S^{\sharp} \subset D
$$

Outcome Semantics


Dependency Semantics


## Dependency Semantics

partitioning with respect to

## 14 <br> 



## Naïve Abstraction

## Naïve Backward Analysis

(2) forget the values of the
sensitive input nodes

(3) check for intersection:
empty $\rightarrow$ fair
otherwise $\rightarrow$ alarm


## Naïve Backward Analysis



## Back to the Semantics...

## Hierarchy of Semantics

parallel semantics


## Parallel Semantics

## 14

 that satisfies dependency fairness with respect to the non-sensitive inputs yields sets of traces that also satisfy dependency fairness

## Parallel Semantics

 that satisfies dependency fairness with respect to the non-sensitive inputs yields sets of traces that also satisfy dependency fairness

## Parallel Semantics



$$
\begin{aligned}
& \alpha_{\square}(S) \stackrel{\text { def }}{=}\left\{\left\{\left\langle t_{0}, t_{\omega}\right\rangle \in R \mid t_{0} \in I\right\} \mid R \in S \wedge I \in \mathbb{Q}\right\} \quad \text { parallel abstraction } \\
& \{[M]\}_{\sim}^{0} \stackrel{\text { def }}{=} \alpha_{0}\left(\llbracket M \rrbracket_{\sim}\right) \\
& =\left\{\left\{\left\langle t_{0}, t_{\omega}\right\rangle \in \Sigma \times \Sigma \mid t \in \llbracket M \rrbracket \wedge t_{0} \in I \wedge t_{\omega} \in O\right\}|I \in \mathbb{}| \wedge O \in \mathbb{O}\right\}
\end{aligned}
$$

## Theorem

Lemma
$M \vDash \mathscr{F}_{i} \Leftrightarrow \forall I \in \mathbb{\square}: \forall A, B \in\{[M]\}_{\sim}^{0}:\left(A_{\omega}^{I} \neq\left.\left. B_{\omega}^{I} \Rightarrow A_{0}^{I}\right|_{\neq i} \cap B_{0}^{I}\right|_{\neq i}=\varnothing\right)$

## Better Abstraction

## Forward and Backward Analysis

(1) partition the space of values of the non-sensitive input nodes






## n caterinaurban / Libra

<> Code
(!) Issues
\$\% Pull requests
(-) Actions
(11) Projects
(1) Security
$\sim$ Insights

|  | master $\checkmark$ \% | $\bigcirc 0$ tags | Go to file | Code - |
| :---: | :---: | :---: | :---: | :---: |
| caterinaurban README |  |  | 9 f 830 db on Aug 8 | (1) 53 commits |
|  |  | RQ5 and RQ6 reproducibility |  | 4 months ago |
|  | .gitignore | RQ1 reproducibility |  | 4 months ago |
| $\square$ | LICENSE | Initial prototype |  | 2 years ago |
| 5 | README.md | RQ5 and RQ6 reproducibility |  | 4 months ago |
|  | README.pdf | README |  | 4 months ago |
|  | icon.png | icon |  | 4 months ago |
|  | libra.png | icon |  | 4 months ago |
|  | requirements.txt | some documentation |  | 4 months ago |
|  | setup.py | some documentation |  | 4 months ago |

Abou

No description or website provided.
\#abstract-interpretation
\#static-analysis
\#machine-learning
\#neural-networks \#fairness

1) Readme
$\Delta \triangle$ MPL-2.0 License

Releases
No releases published

## Packages

No packages published

## Languages

- Python $98.7 \%$
- Shell $1.3 \%$

Nowadays, machine-learned software plays an increasingly important role in critical decision-making in our social, economic, and civic lives.

## Formal Methods for Model Training

## Robust Training

## Minimizing the Worst-Case Loss for Each Input



## Robust Training

## Minimizing the Worst-Case Loss for Each Input

## Adversarial Training

Minimizing a Lower Bound on the Worst-Case Loss for Each Input

$$
\max _{\mathbf{x}^{\prime} \in \mathcal{C}(\mathbf{x})} \mathcal{L}\left(f\left(\boldsymbol{\theta}, \mathbf{x}^{\prime}\right), \mathbf{y}\right)
$$



$$
\begin{gathered}
\mathrm{VI} \\
\mathcal{L}\left(f\left(\boldsymbol{\theta}, \mathbf{x}_{\text {adv }}\right), y\right)
\end{gathered}
$$

generate adversarial inputs and use them as training data

# Robust Training <br> <br> Minimizing the Worst-Case Loss for Each Input 

 <br> <br> Minimizing the Worst-Case Loss for Each Input}

## Certified Training

$$
\max _{\mathbf{x}^{\prime} \in \mathcal{C}(\mathbf{x})} \mathcal{L}\left(f\left(\boldsymbol{\theta}, \mathbf{x}^{\prime}\right), \mathbf{y}\right)
$$

Minimizing an Upper Bound on the Worst-Case Loss for Each Input


$$
\mathcal{L}_{\text {ver }}(f(\boldsymbol{\theta}, \mathbf{x}), y)
$$

use upper bound as regularizer to encourage robustness

## Robust Training

## Minimizing the Worst-Case Loss for Each Input

## Hybrid Training

## Minimizing an Approximation of the Worst-Case Loss

 that Contains and Adversarial Example for Each Input

## Bibliography

[Kurd03] Zeshan Kurd, Tim Kelly. Establishing Safety Criteria for Artificial Neural Networks. In KES, pages 63-169, 2003.
[Li19] 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 SAS, page 296-319, 2019.
[Singh19] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin T. Vechev. An Abstract Domain for Certifying Neural Networks. In POPL, pages 41:1-41:30, 2019.
[Mazzucato21] Denis Mazzucato and Caterina Urban. Reduced Products of Abstract Domains for Fairness Certification of Neural Networks. In SAS, 2021.
[Julian16] Kyle D. Julian, Jessica Lopez, Jeffrey S. Brush, Michael P. Owen, Mykel J. Kochenderfer. Policy Compression for Aircraft Collision Avoidance Systems. In DASC, pages 1-10, 2016.
[Katz17] Guy Katz, Clark W. Barrett, David L. Dill, Kyle Julian, Mykel J. Kochenderfer. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks. In CAV, pages 97-117, 2017.
[Galhotra17] Sainyam Galhotra, Yuriy Brun, and Alexandra Meliou.
Fairness Testing: Testing Software for Discrimination. In FSE, pages 498-510, 2017.
[Urban20] Caterina Urban, Maria Christakis, Valentin Wüstholz, and
Fuyuan Zhang. Perfectly Parallel Fairness Certification of Neural Networks. In OOPSLA, pages 185:1-185:30, 2020.
[Urban21] Caterina Urban and Antoine Miné. A Review of Formal Methods applied to Machine Learning. https://arxiv.org/abs/2104.02466, 2021.
[Müller23] Mark Niklas Müller, Franziska Eckert, Marc Fischer, Martin Vechev. Certified Training: Small Boxes Are All You Need. In ICLR, 2023.


[^0]:    R. Ehlers - Formal Verification of Piece-Wise Linear Feed-Forward Neural Networks (ATVA 2017)

[^1]:    S. Dutta et al. - Output Range Analysis for Deep Feedforward Neural Networks (NFM 2018)

[^2]:    S. Dutta et al. - Output Range Analysis for Deep Feedforward Neural Networks (NFM 2018)

[^3]:    H.-D. Tran et al. - Star-Based Reachability Analysis of Deep Neural Networks (FM 2018)

[^4]:    P. Prabhakar and Z. R. Afza - Abstraction based Output Range Analysis for Neural Networks (NeurIPS 2019)

