FDL 2020

Robust Models for Source Code: Techniques and Applications

Pavol Bielik
pavol.bielik@inf.ethz.ch

Department of Computer Science
ETH zürich

SRILAB

Secure, Reliable, and Intelligent Systems Lab @ ETH

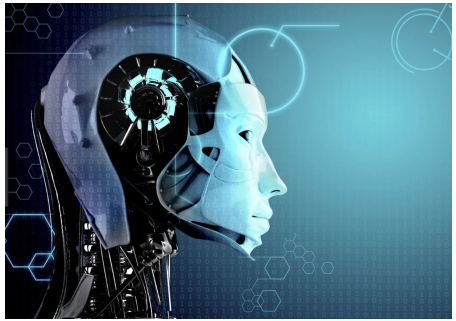

Robust ML

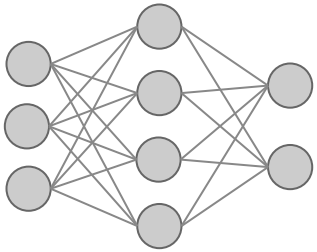

Neural network verification

ML for programming

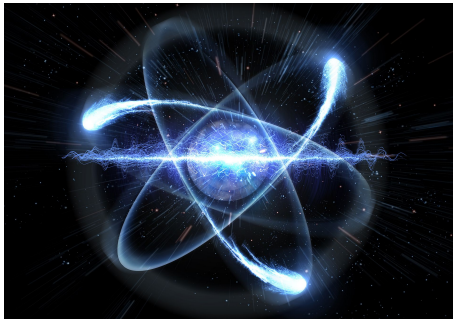

Probabilistic programming

Secure, Reliable, and Intelligent Systems Lab @ ETH

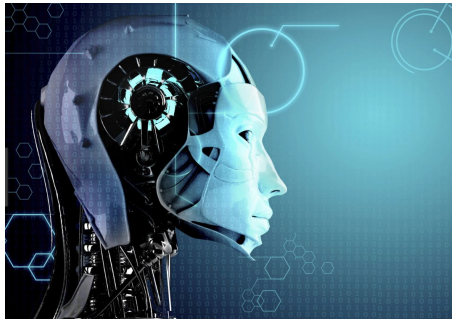

Robust ML

Neural network verification

ML for programming

Probabilistic programming

Statistical Programming Tools

Code Completion

```
Camera camera = Camera.open();
camera.SetDisplayOrientation(90);
                ?
```

Port Code

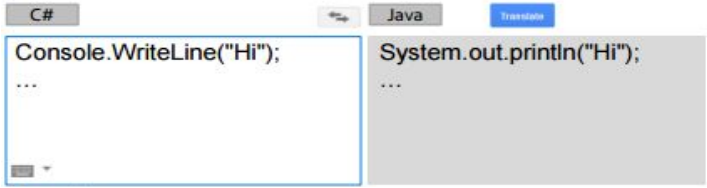

Program Synthesis

Learning to Solve Formulas

φ → **Z3** → SAT + model $\mathcal{M} \models \varphi$
SMT Solver UNSAT + unsat core

(check-sat-using </> Strategy)

Up to 100x speed-up over Z3

Bug Detection

likely error

```
...
for x in range(a):
    print a[x]
```

Deep Learning + Code

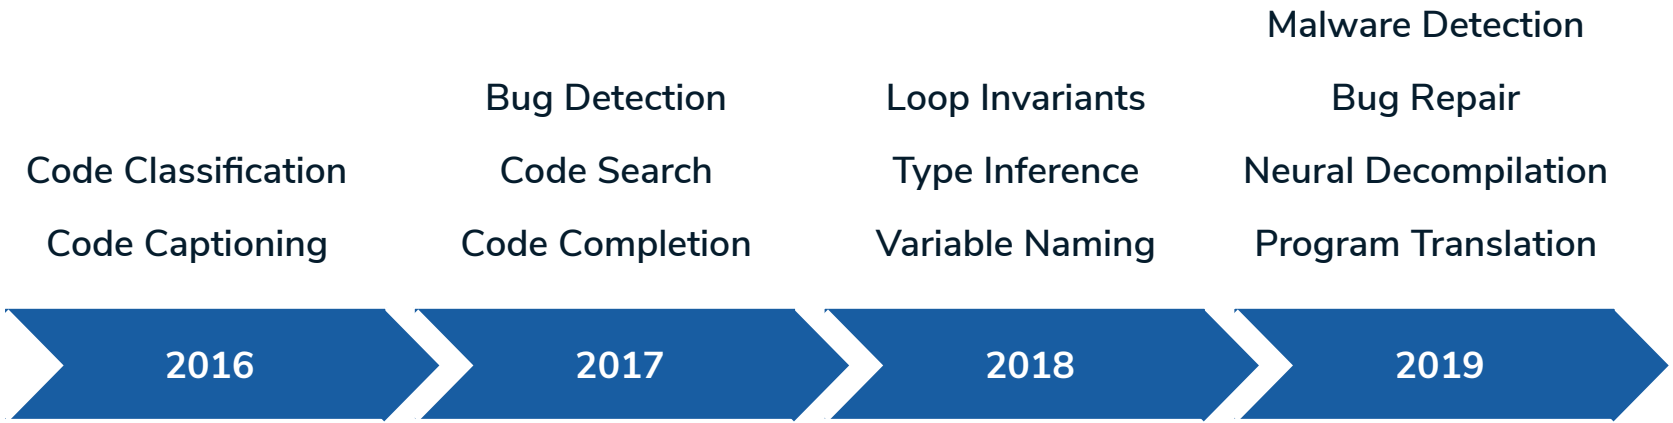

Deep Learning + Code

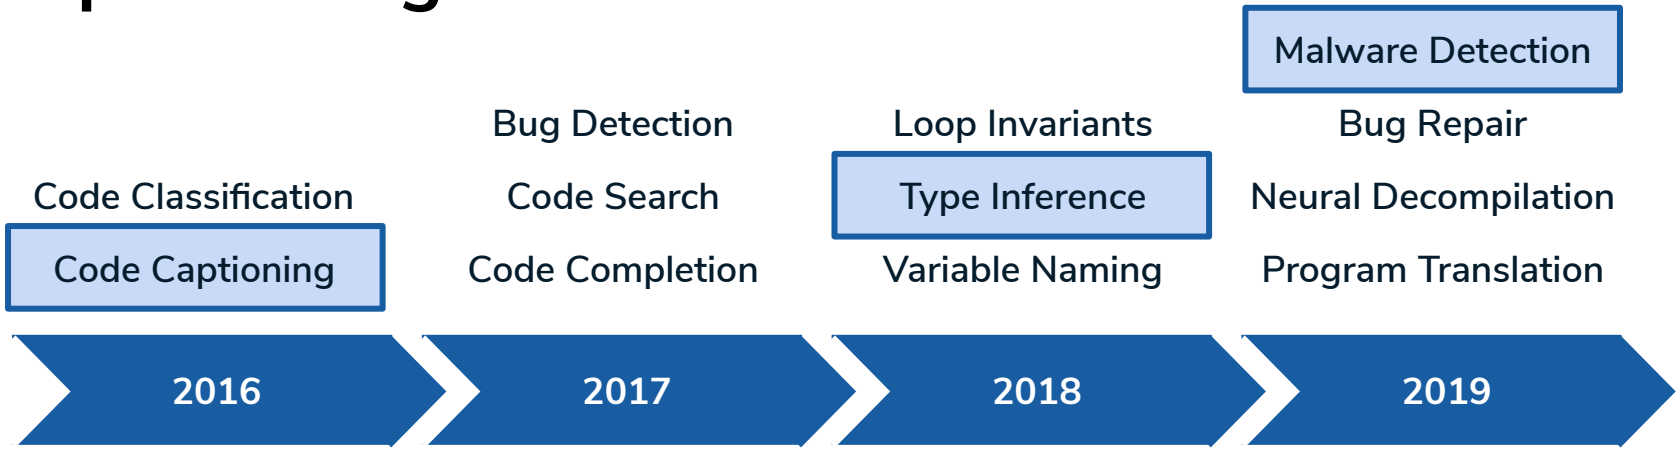

Techniques are general and apply to other tasks

Deep Learning + Code

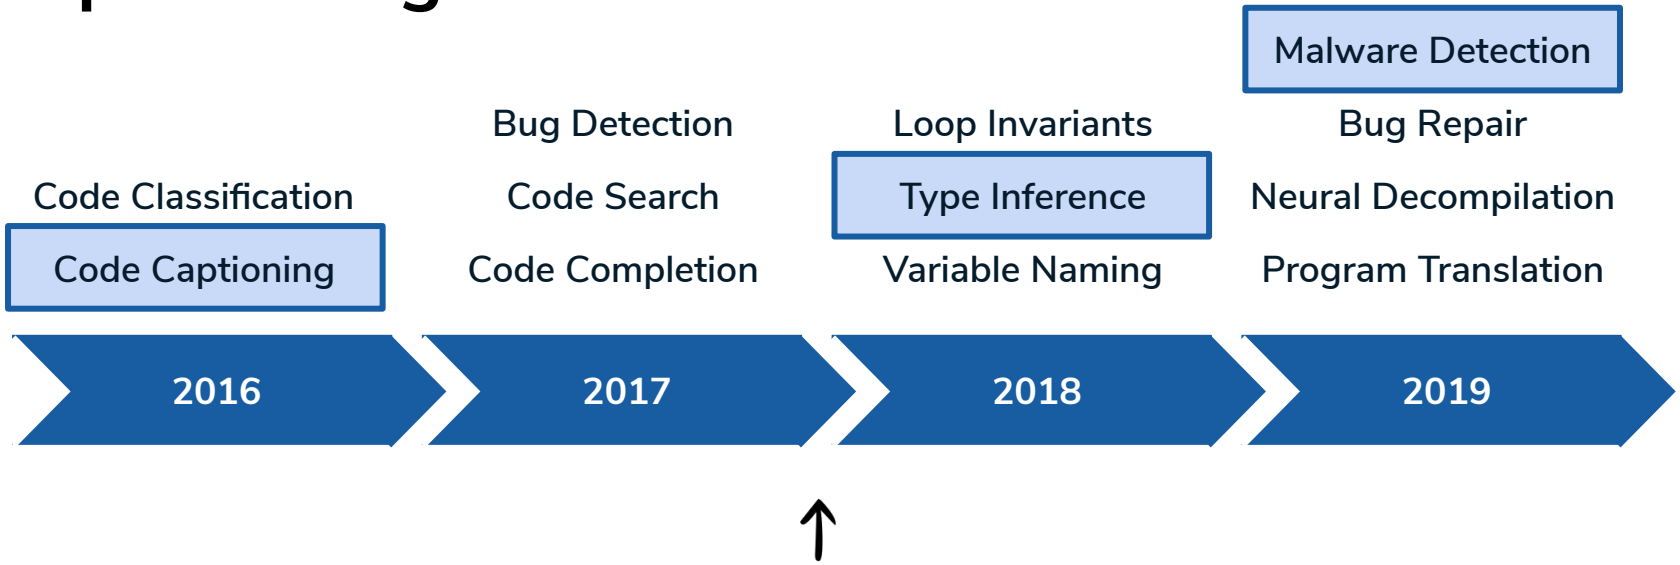

Majority is based on deep learning models

Techniques are general and apply to other tasks

Secure, Reliable, and Intelligent Systems Lab @ ETH

Robust ML

Neural Network Verification

ML for programming

Probabilistic programming

Adversarial Robustness

Vision

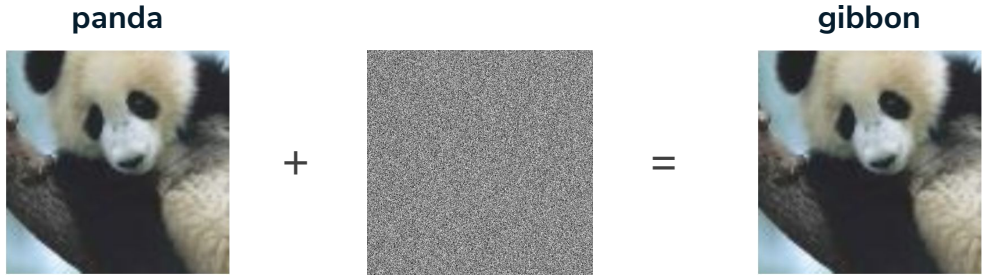

Explaining and Harnessing Adversarial Examples. Goodfellow et. al. ICLR'15

Adversarial Robustness

Vision

stop sign

+

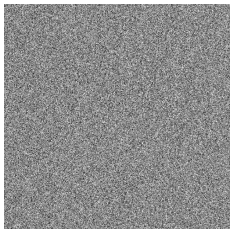

=

gibbon

Adversarial Robustness

Vision

stop sign

+

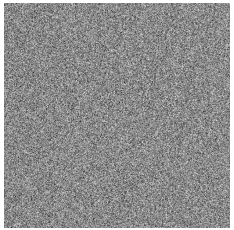

=

gibbon

Sound

+ **noise** =

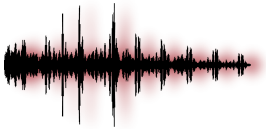

Audio Adversarial Examples: Targeted Attacks on Speech-to-Text. Carlini et. al. ICML'18 workshop

Adversarial Robustness

Vision

stop sign

+

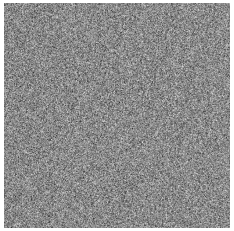

=

gibbon

Sound

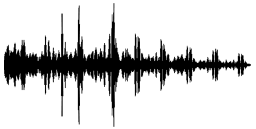

+ **noise** =

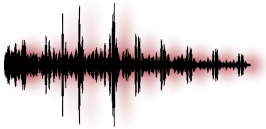

Audio Adversarial Examples: Targeted Attacks on Speech-to-Text. Carlini et. al. ICML'18 workshop

Code

+ **code refactoring** =

Adversarial Robustness

stop sign

gibbon

Vision

+

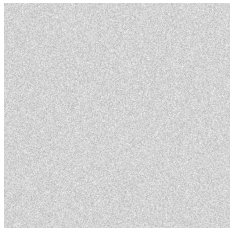

=

Sound

+

noise

=

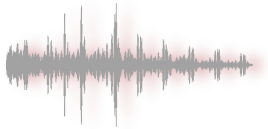

Audio Adversarial Examples: Targeted Attacks on Speech-to-Text. Carlini et. al. ICML'18 workshop

Code

+

code refactoring

=

Adversarial Robustness for Code

1 **How robust are existing models?**

Adversarial Robustness for Code

1 **How robust are existing models?**

Adversarial Robustness for Code

1. **How robust are existing models?**
2. **How to find adversarial examples?**

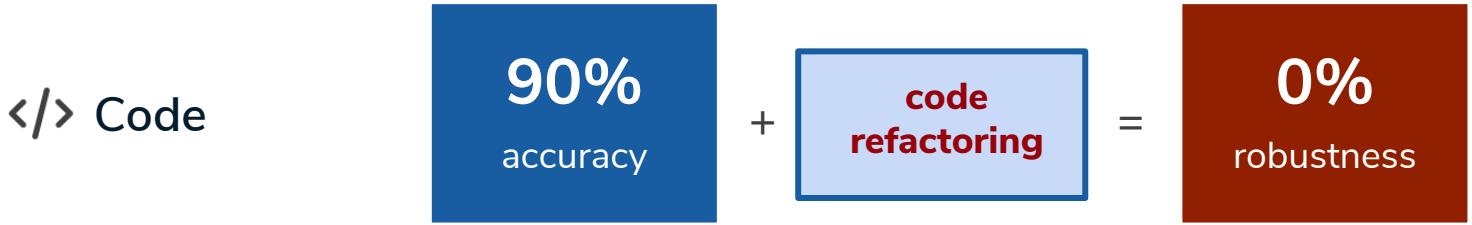

Adversarial Robustness for Code

1. **How robust are existing models?**
2. **How to find adversarial examples?**
3. **How to improve robustness?**

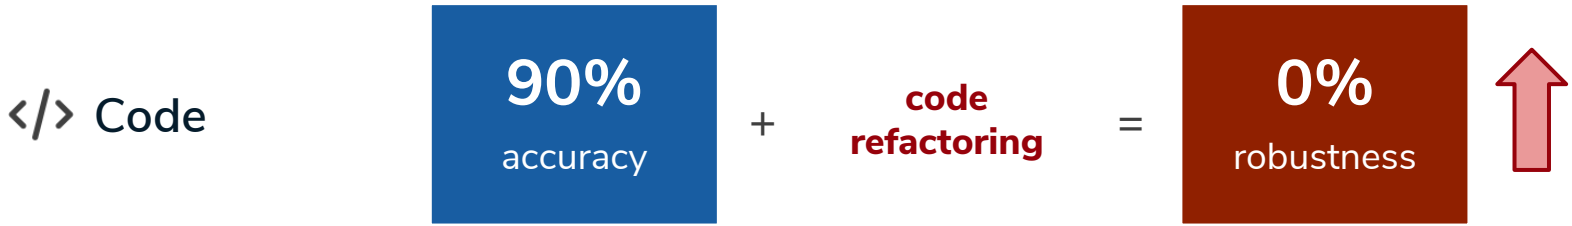

Adversarial Robustness Example

Input Program
x

```
...
v = parseInt(
  hex.substr(1),
  radix
)
...
```

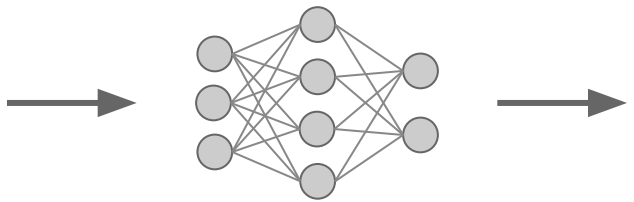

(Type Inference)
Program Properties
y

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

Adversarial Robustness Example

Input Program
x

```
...
v = parseInt(
  hex.substr(1),
  radix
)
...
```

Model
$f(\mathbf{x}) \rightarrow \mathbf{y}$

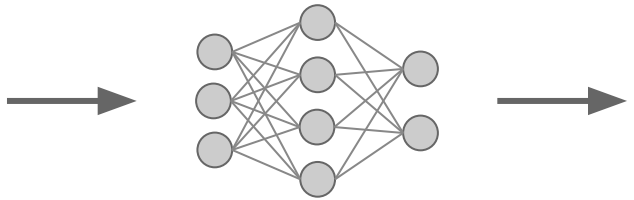

(Type Inference)
Program Properties
y

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

Goal (Adversarially Robustness):
Model is correct for *all* label preserving program transformations

```
...
v = parseInt(
  color.substr(1),
  radix
)
...
```

variable renaming

```
...
v = parseInt(
  hex.substr(42),
  radix
)
...
```

constant replacement

```
...
v = parseInt(
  hex.substr(1),
  radix + 0
)
...
```

semantic equivalence

```
...
parseInt(
  hex.substr(1),
  radix
)
...
```

remove assignment

Adversarial Robustness Example

Input Program
x

```
...
v = parseInt(
  hex.substr(1),
  radix
)
...
```

Model
$f(\mathbf{x}) \rightarrow \mathbf{y}$

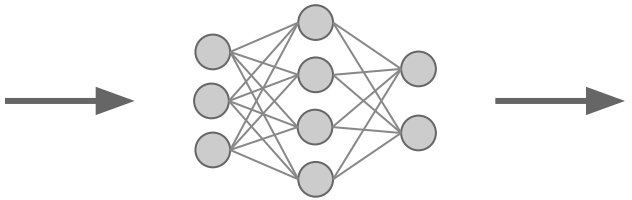

(Type Inference)
Program Properties
y

```
...
v[num] = parseInt[num](
  hex[str].substr[str](1),
  radix[num]
)
...
```

Goal (Adversarially Robustness):
Model is correct for *all* label preserving program transformations

$S(x)$

Set of valid program transformations for **x**

```
...
v = parseInt(
  color.substr(1),
  radix
)
...
```

variable renaming

```
...
v = parseInt(
  hex.substr(42),
  radix
)
...
```

constant replacement

```
...
v = parseInt(
  hex.substr(1),
  radix + 0
)
...
```

semantic equivalence

```
...
parseInt(
  hex.substr(1),
  radix
)
...
```

remove assignment

How robust are existing models?

Type Inference [1]

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

[1] Adversarial Robustness for Code. ICML'20

How robust are existing models?

Type Inference [1]

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  int i = 0;
  for (Object elem: this.elements) {
    if (elem.equals(target))
      return i
    i++;
  return -1;
}
```

[1] Adversarial Robustness for Code. ICML'20
[2] Adversarial Examples for Models of Code. ArXiv'19
[3] Semantic Robustness of Models of Source Code. ArXiv'20

How robust are existing models?

Type Inference [1]

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  int i = 0;
  for (Object elem: this.elements) {
    if (elem.equals(target))
      return i
    i++;
  return -1;
}
```

Malware Detection [4]

```
push ebp
mov ebp, esp
push ebx
push edx                    yes
mov ebx, [ebp+4]            no
add ebx, 0x10
mov edx, [ebp+8]
mov [edx], ebx
```


[1] Adversarial Robustness for Code. ICML'20
[2] Adversarial Examples for Models of Code. ArXiv'19
[3] Semantic Robustness of Models of Source Code. ArXiv'20
[4] Optimization-Guided Binary Diversification to Mislead Neural Networks for Malware Detection. ArXiv'19

How robust are existing models?

Type Inference [1]

```
...
v = parseInt(
  hex.substr(1),
  radix
)
...
```

89% accuracy

48% robustness

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  int i = 0;
  for (Object elem: this.elements) {
    if (elem.equals(target))
      return i
    i++;
  return -1;
}
```

39.2 F1

19.6 robust F1

Malware Detection [4]

```
push ebp
mov ebp, esp
push ebx
push edx                yes
mov ebx, [ebp+4]        no
add ebx, 0x10
mov edx, [ebp+8]
mov [edx], ebx
```

99% accuracy

~2% robustness

[1] Adversarial Robustness for Code. ICML'20
[2] Adversarial Examples for Models of Code. ArXiv'19
[3] Semantic Robustness of Models of Source Code. ArXiv'20
[4] Optimization-Guided Binary Diversification to Mislead Neural Networks for Malware Detection. ArXiv'19

How robust are existing models?

Type Inference [1]

```
...
```

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  return -1;
}
```

Malware Detection [4]

```
push ebp
mov edx, [ebp+8]
mov [edx], ebx
```

Not trained with robustness in mind

89% accuracy	48% robustness

39.2 F1	19.6 robust F1

99% accuracy	~2% robustness

[1] Adversarial Robustness for Code. ICML'20
[2] Adversarial Examples for Models of Code. ArXiv'19
[3] Semantic Robustness of Models of Source Code. ArXiv'20
[4] Optimization-Guided Binary Diversification to Mislead Neural Networks for Malware Detection. ArXiv'19

So Far...

original programs

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

δ: rename
hex ⟶ color

transformed programs

```
...
vnum = parseIntnum(

colornum.substrnum(1),
  radixnum
)
...
```

89%
accuracy

48%
robustness

Adversarial Training

original programs

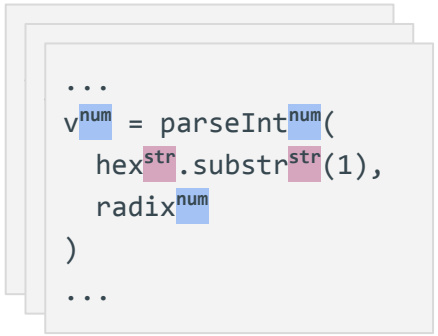

δ: rename
hex ⟶ color

transformed programs

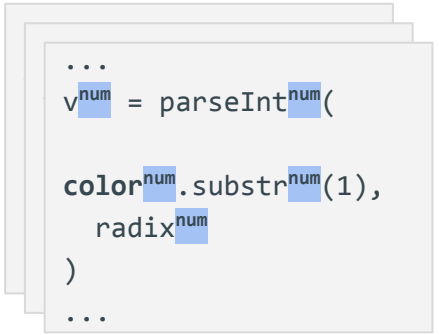

Standard training

$$\min loss(\theta, \mathbf{x}, \mathbf{y})$$

measures the model performance

Adversarial training

$$\min [\max loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})]$$
$$\boldsymbol{\delta} \in \mathbf{S(x)}$$

program transformations

Adversarial Training

original programs

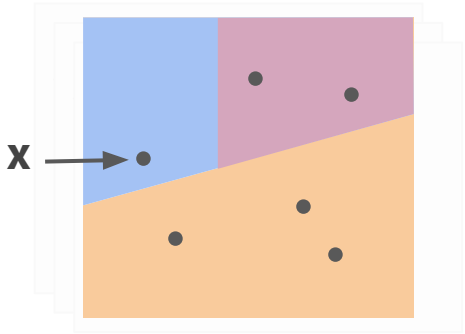

δ: rename
hex ⟶ color

generated programs

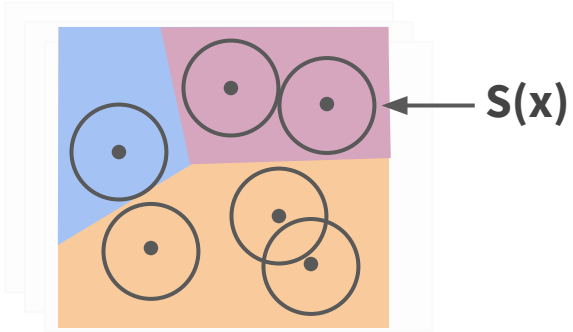

Standard training

$\min loss(\theta, \mathbf{x}, \mathbf{y})$

trains on individual programs

Adversarial training

$\min [\max loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})]$
$\boldsymbol{\delta} \in \mathbf{S(x)}$

trains on worst case generated programs

Finding Adversarial Examples

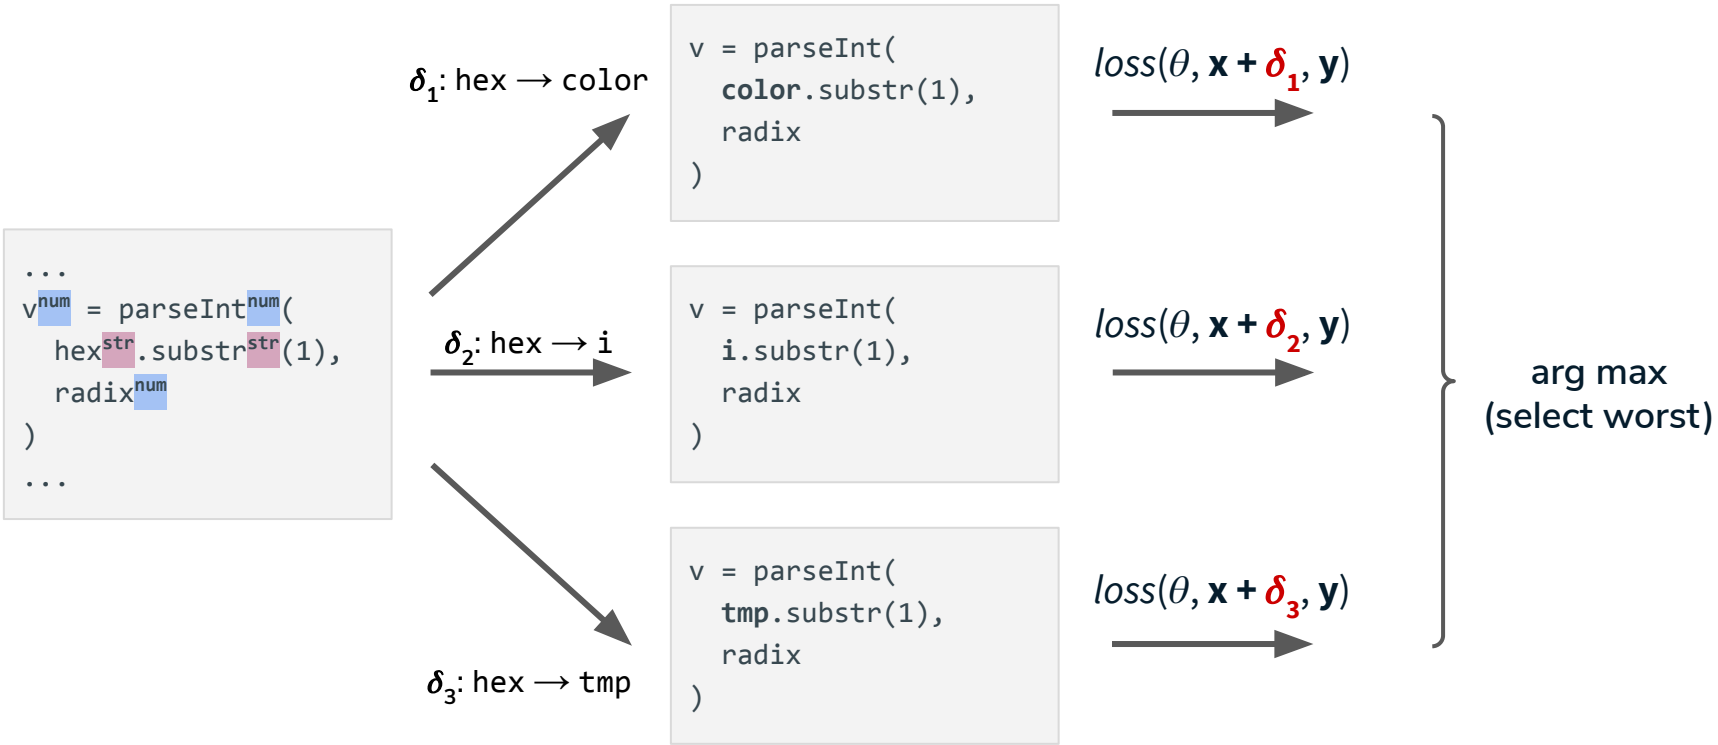

Basic approach: Try all valid modifications

Word Embeddings

maps each **discrete word** to a **continuous vector**

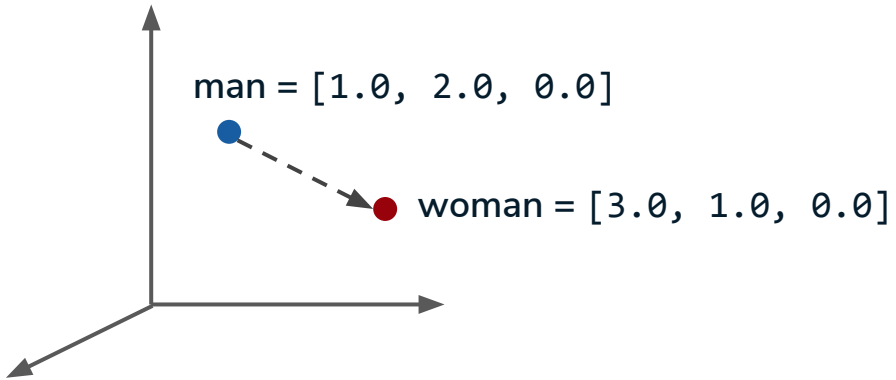

Distributed Representations of Words and Phrases and their Compositionality. Mikolov et. al. NeurIPS'13

Word Embeddings

maps each **discrete word** to a **continuous vector**

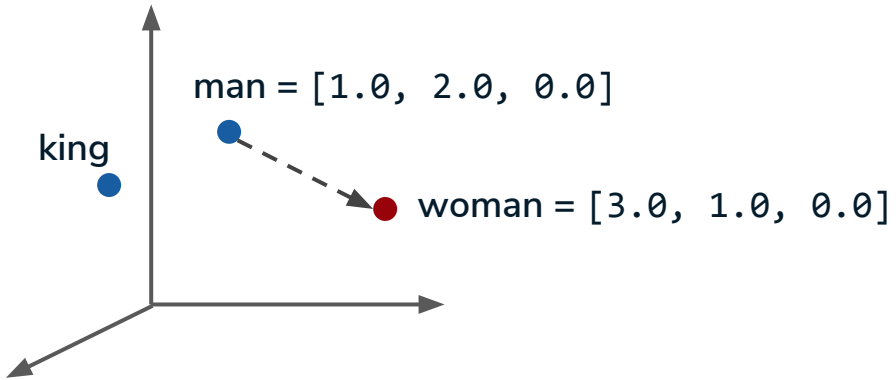

Distributed Representations of Words and Phrases and their Compositionality. Mikolov et. al. NeurIPS'13

Word Embeddings

maps each **discrete word** to a **continuous vector**

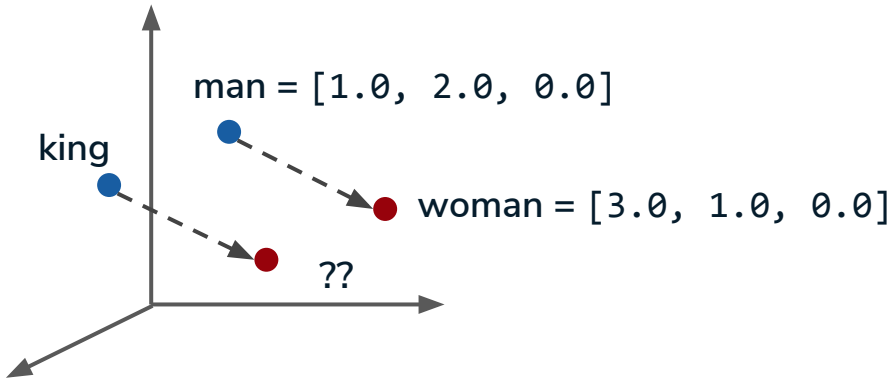

Distributed Representations of Words and Phrases and their Compositionality. Mikolov et. al. NeurIPS'13

Word Embeddings

maps each **discrete word** to a **continuous vector**

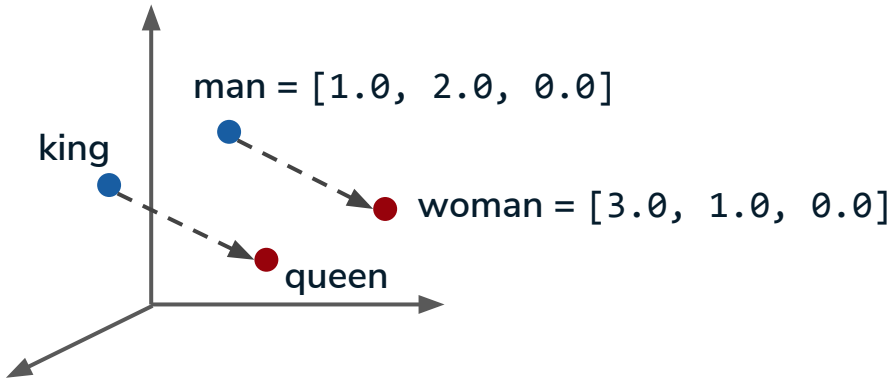

Distributed Representations of Words and Phrases and their Compositionality. Mikolov et. al. NeurIPS'13

Finding Adversarial Examples

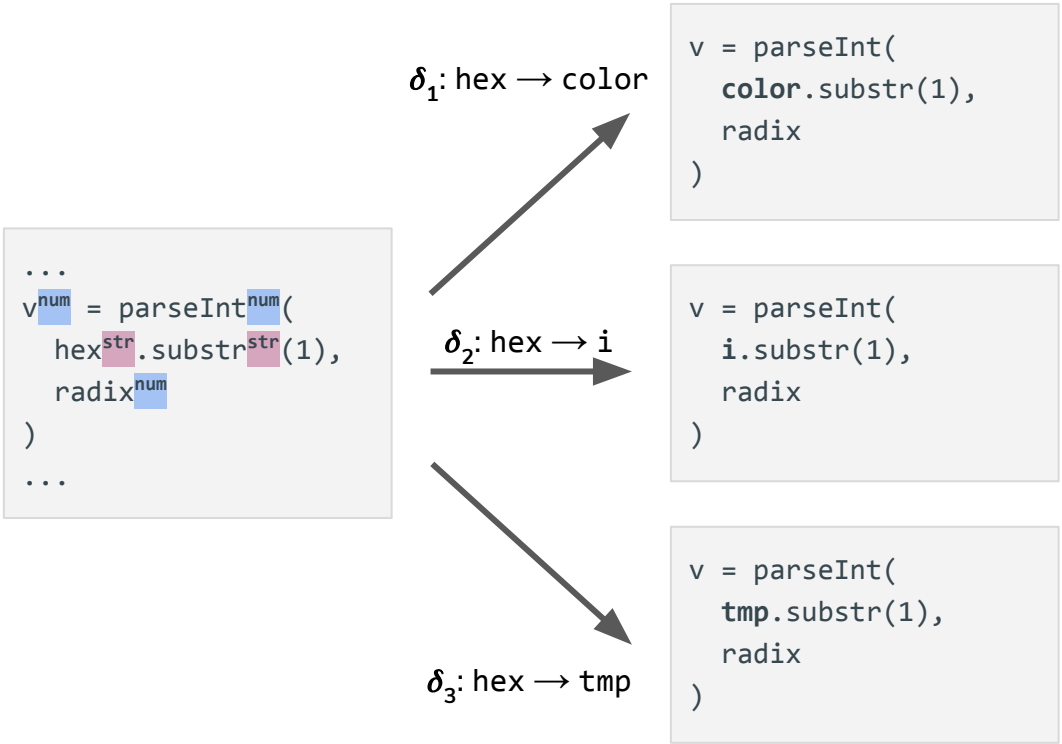

Finding Adversarial Examples

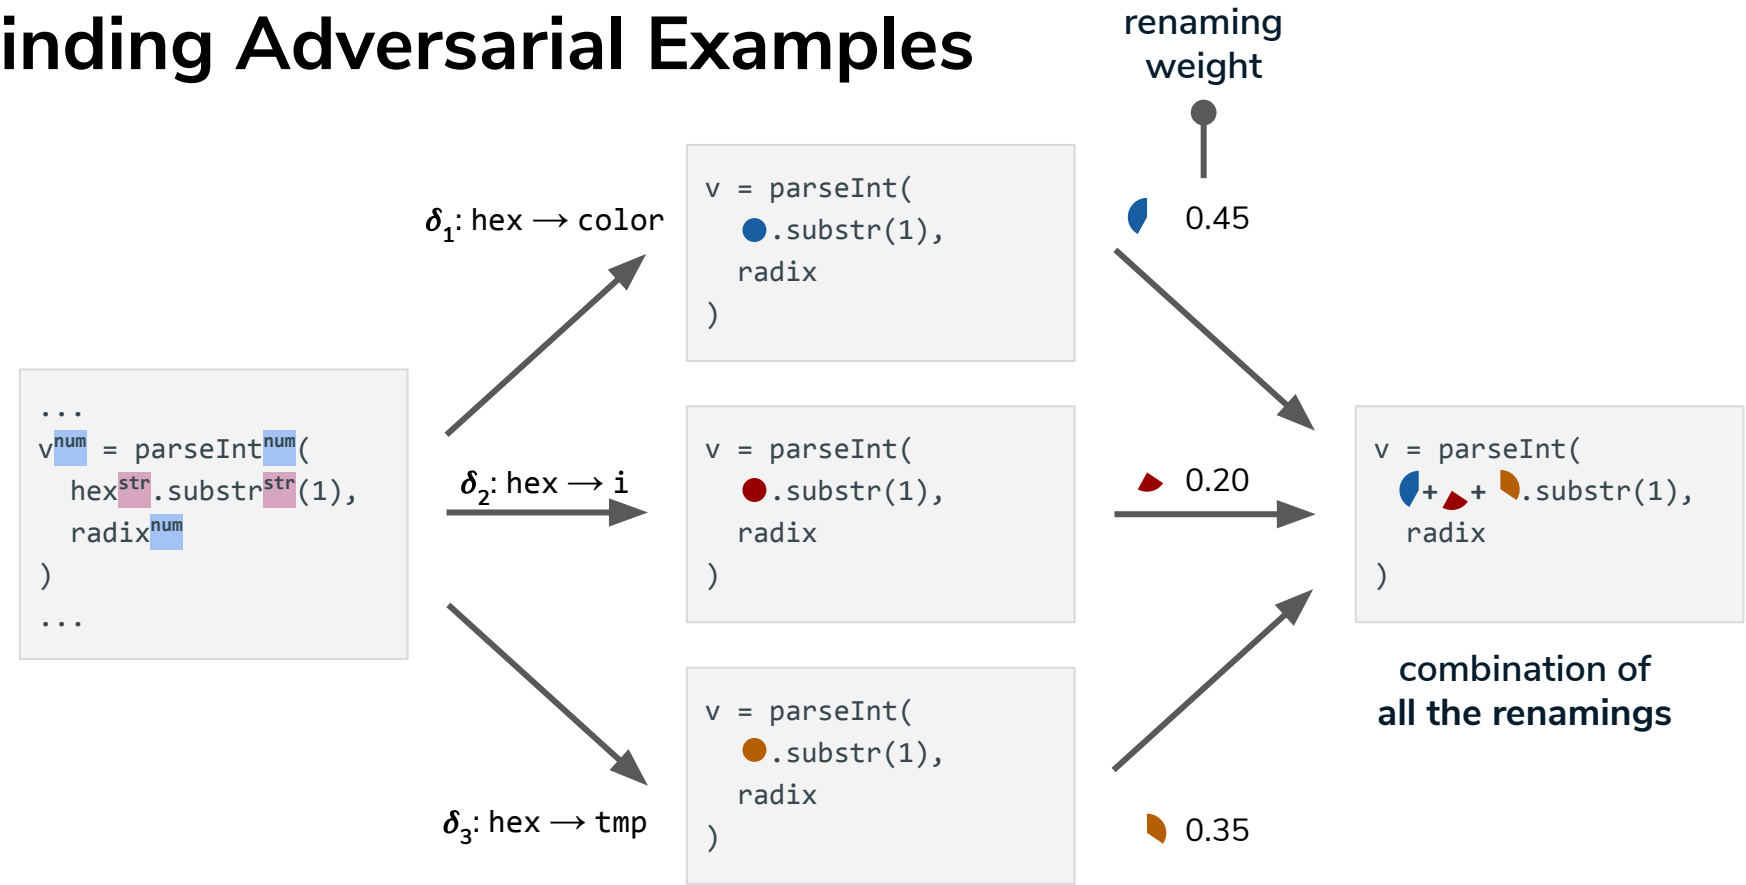

Finding Adversarial Examples

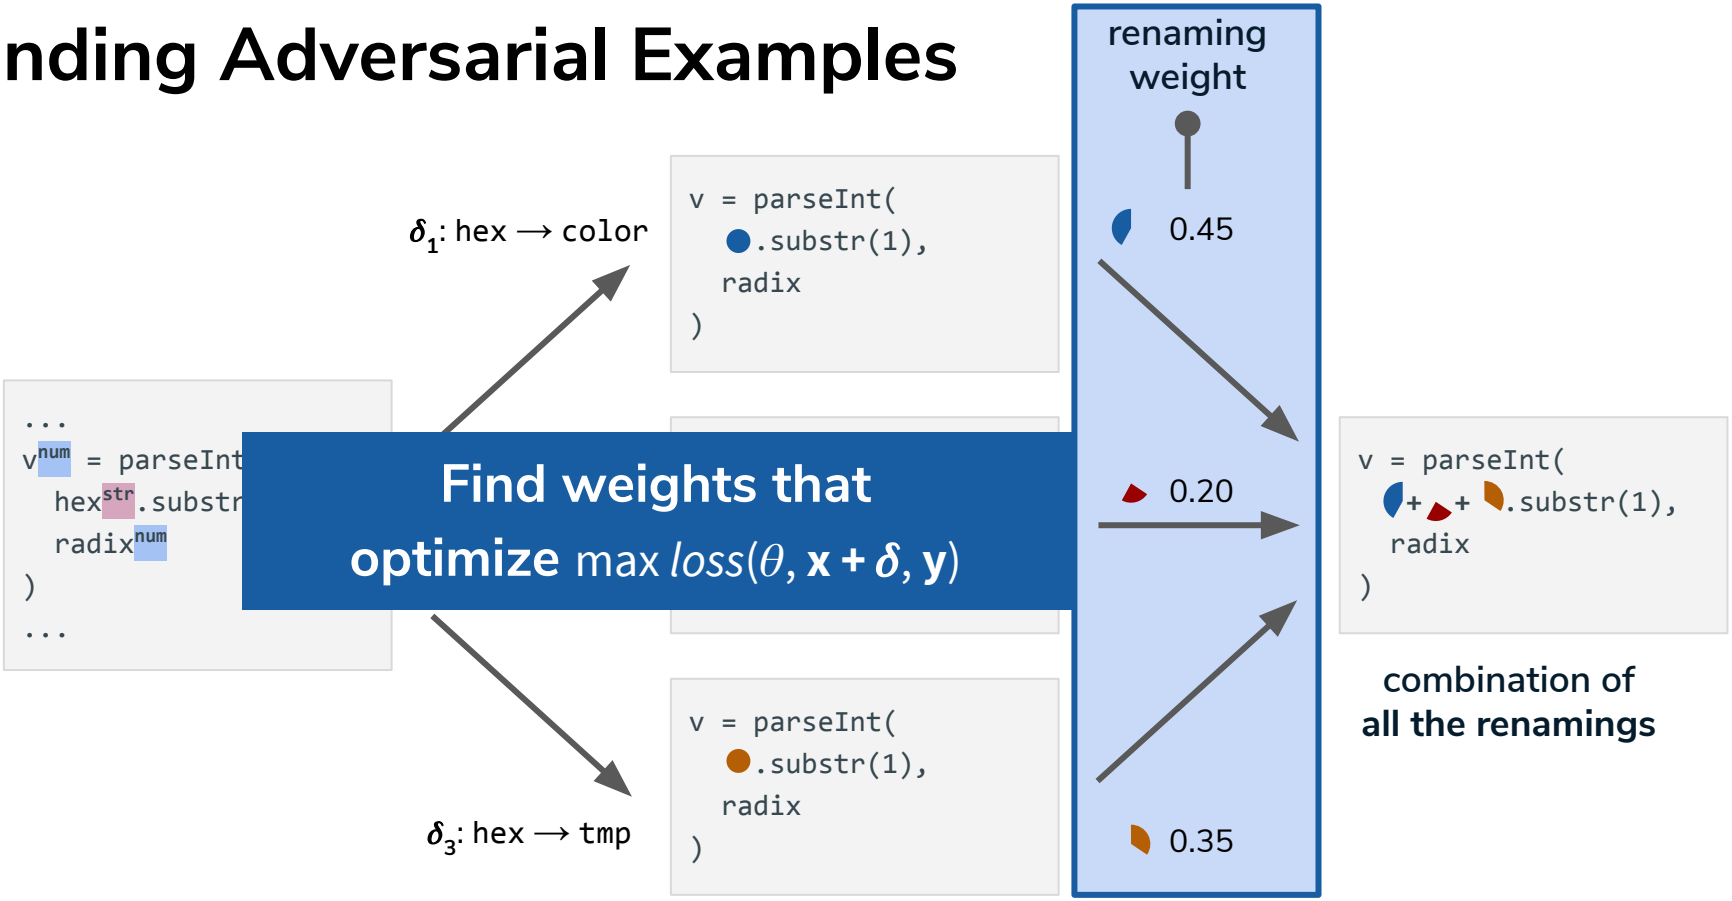

Solving the Inner max *loss* Efficiently

Gradient Based Optimization

$$\theta \leftarrow \theta - \nabla_{\delta} loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y}) \quad \boldsymbol{\delta} \in \mathbf{S(x)}$$

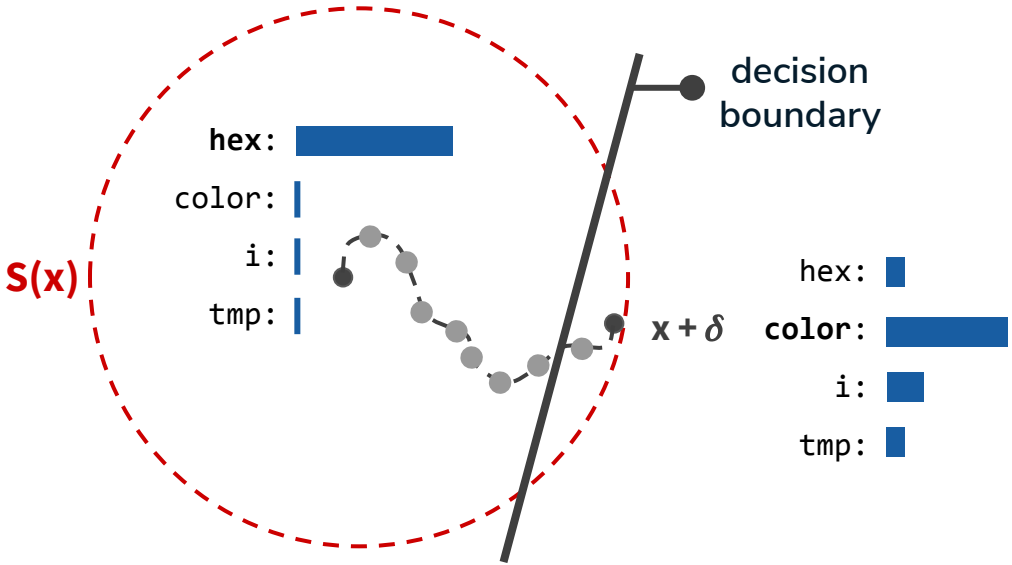

Adversarial Examples for Models of Code.
Yefet et. al. ArXiv'20

Robustness of Existing Models

Type Inference [1]

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

89% accuracy

45% robustness

Adversarial training

57% robustness

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  int i = 0;
  for (Object elem: this.elements) {
    if (elem.equals(target))
      return i
    i++;
  return -1;
}
```

39.2 F1

19.6 robust F1

Malware Detection [4]

```
push ebp
mov ebp, esp
push ebx
push edx                    yes
mov ebx, [ebp+4]            no
add ebx, 0x10
mov edx, [ebp+8]
mov [edx], ebx
```

99% accuracy

~2% robustness

Robustness of Existing Models

Type Inference [1]

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

89%	45%
accuracy	robustness
Adversarial training	57% robustness

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  int i = 0;
  for (Object elem: this.elements) {
    if (elem.equals(target))
      return i
    i++;
  return -1;
}
```

39.2	19.6
F1	robust F1
40.6	27.5
F1	robust F1

Malware Detection [4]

```
push ebp
mov ebp, esp
push ebx
push edx                    yes
mov ebx, [ebp+4]            no
add ebx, 0x10
mov edx, [ebp+8]
mov [edx], ebx
```

99%	~2%
accuracy	robustness

Robustness of Existing Models

Type Inference [1]

```
...
vnum = parseIntnum(
  hexstr.substrstr(1),
  radixnum
)
...
```

89% accuracy	45% robustness
Adversarial training	57% robustness

Code Captioning [2,3]

```
int indexOfTarget(Object target) {
  int i = 0;
  for (Object elem: this.elements) {
    if (elem.equals(target))
      return i
    i++;
  return -1;
}
```

39.2 F1	19.6 robust F1
40.6 F1	27.5 robust F1

Malware Detection [4]

```
push ebp
mov ebp, esp
push ebx
push edx                    yes
mov ebx, [ebp+4]            no
add ebx, 0x10
mov edx, [ebp+8]
mov [edx], ebx
```

99% accuracy	~2% robustness
too expensive	

Solving the Inner max *loss* Efficiently

Gradient Based Optimization

$$\theta \leftarrow \theta - \nabla_{\delta} loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})$$
$$\boldsymbol{\delta} \in \mathbf{S(x)}$$

```
...
v = parseInt(
  color.substr(1),
  radix
)
...
```

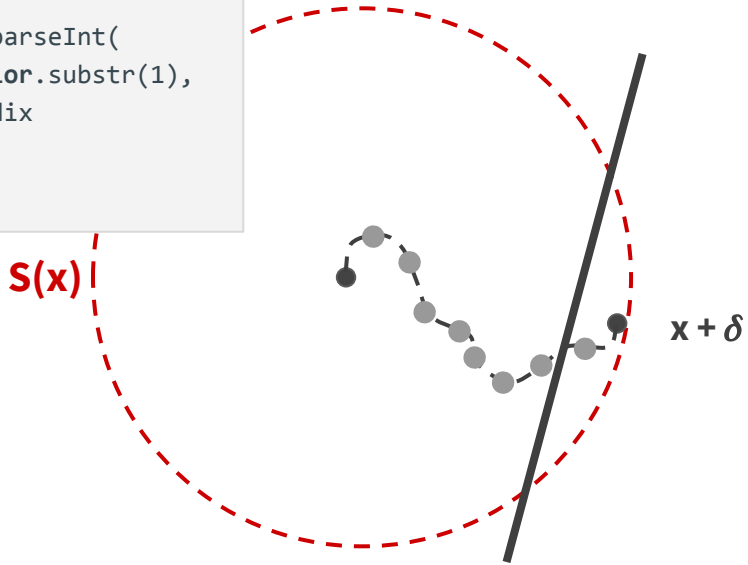

Refine S

$$\min [\max loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})]$$
$$\boldsymbol{\delta} \in \mathbf{S}(\alpha(\mathbf{x}))$$

```
parseInt(
  _,
  _
)
```

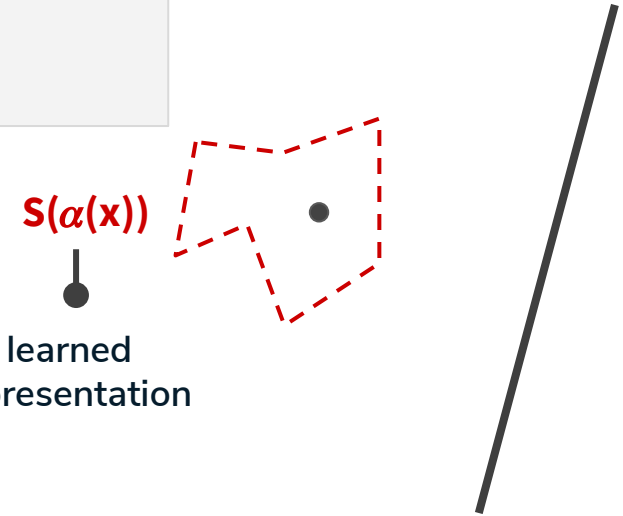

Adversarial Robustness for Code.
Bielik et. al. ICML'20

Solving the Inner max *loss* Efficiently

Gradient Based Optimization

$$\theta \leftarrow \theta - \nabla_{\boldsymbol{\delta}} loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})$$

$$\boldsymbol{\delta} \in \mathbf{S(x)}$$

```
...
v = parseInt(
  color.substr(1),
  radix
)
...
```

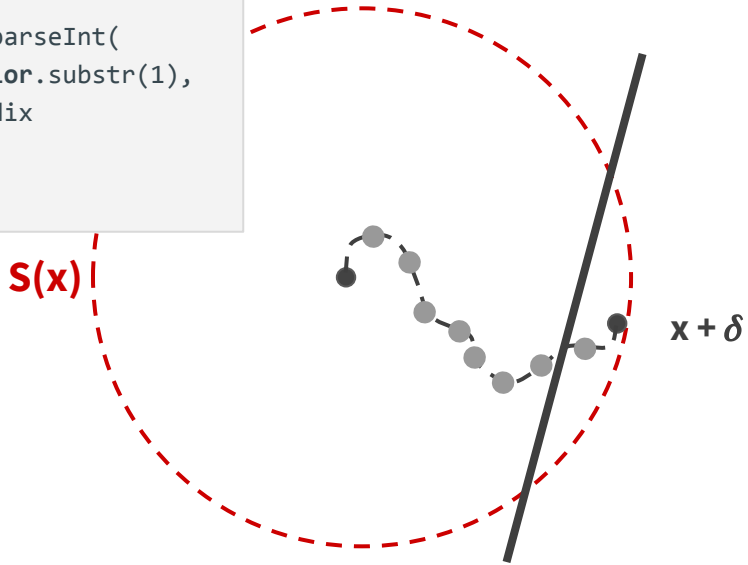

Refine S

$$\min [\max loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})]$$

$$\boldsymbol{\delta} \in \mathbf{S}(\alpha(\mathbf{x}))$$

```
parseInt(
  _,
  _
)
```

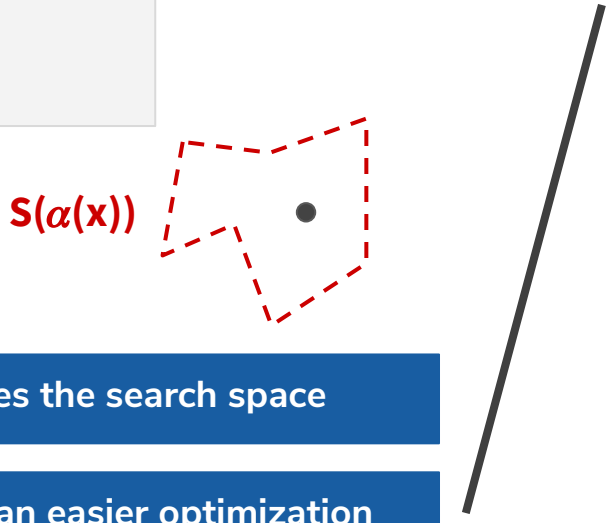

reduces the search space

leads to an easier optimization

Solving the Inner max *loss* Efficiently

Gradient Based Optimization

$$\theta \leftarrow \theta - \nabla_{\delta} loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})$$
$$\boldsymbol{\delta} \in \mathbf{S(x)}$$

```
...
v = parseInt(
  color.substr(1),
  radix
)
...
```

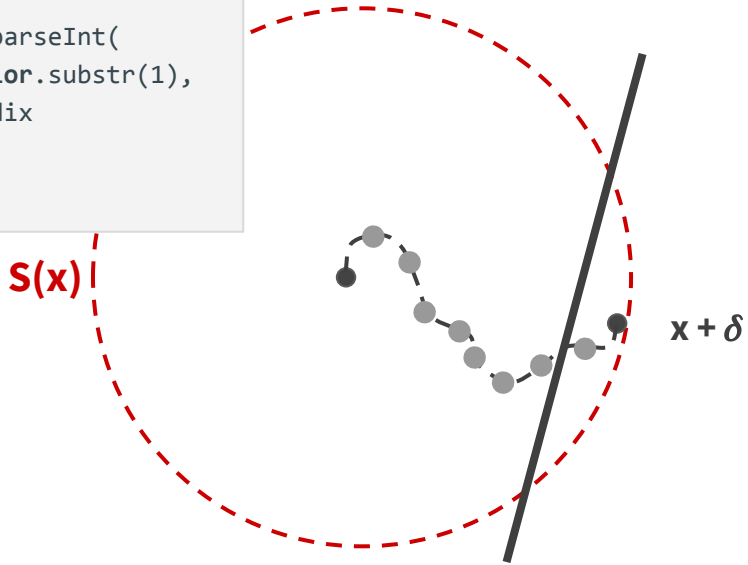

Refine S

$$\min [\max loss(\theta, \mathbf{x} + \boldsymbol{\delta}, \mathbf{y})]$$
$$\boldsymbol{\delta} \in \mathbf{S}(\alpha(\mathbf{x}))$$

- orthogonal to gradient optimization
- supports all transformations
- reduces the search space
- leads to an easier optimization

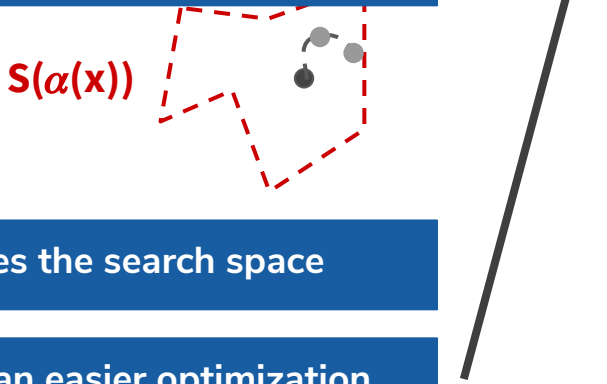

Key Techniques

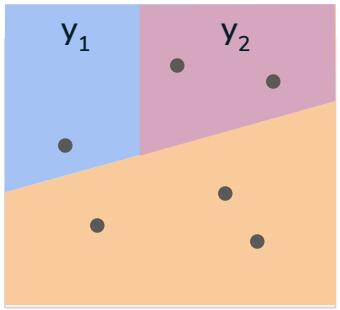

δ = hex ⟶ color

$\delta \in S(x)$

1st technique

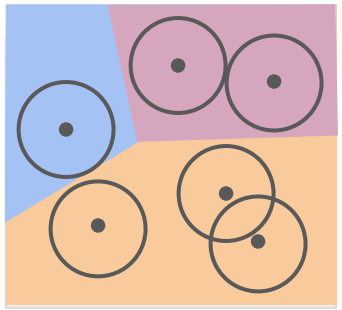

$\alpha(\mathbf{x} + \delta)$

2nd technique

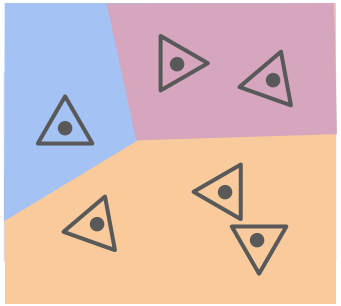

Standard Training

Adversarial Training

Representation Learning

Key Techniques

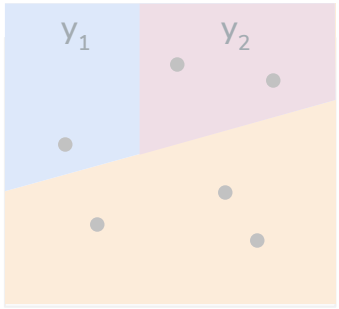

Standard Training

δ = hex ⟶ color

$\delta \in S(x)$

1st technique

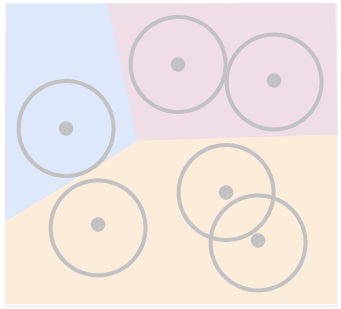

Adversarial Training

$\alpha(\mathbf{x} + \delta)$

2nd technique

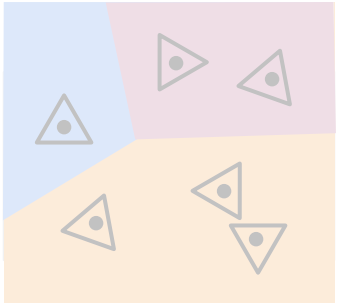

Representation Learning

3rd technique

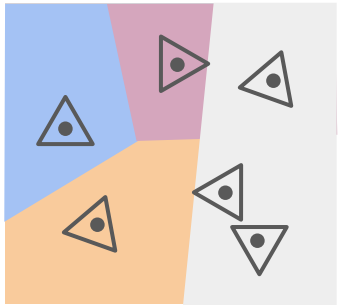

Abstain

=

Predict label

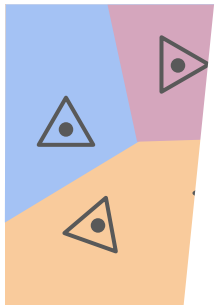

Robust & Accurate

+

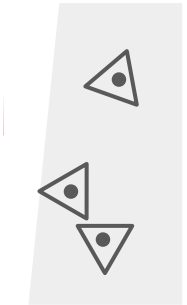

Robust

Not enough confidence for prediction

Robustness of Existing Models

Type Inference [1]		Code Captioning [2,3]		Malware Detection [4]	
89% accuracy	45% robustness	39.2 F1	19.6 robust F1	99% accuracy	~2% robustness
Adversarial training	57% robustness		27.5 robust F1		too expensive
+ Refinement & Abstain	67% robustness				

Robust Models for Source Code

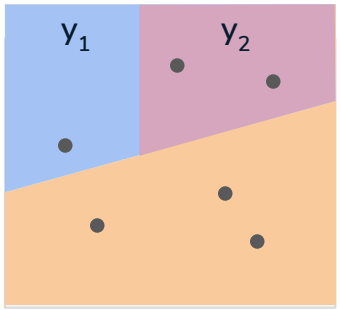

Standard Training

δ = hex ⟶ color

$\delta \in S(x)$

1st technique

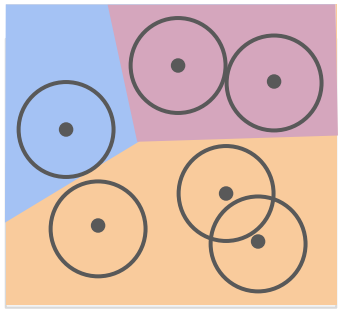

Adversarial Training

$\alpha(\mathbf{x} + \delta)$

2nd technique

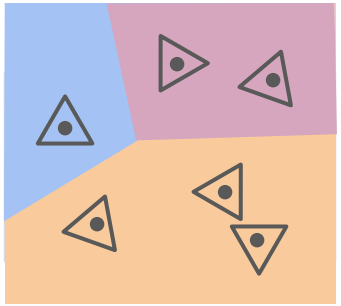

Representation Learning

3rd technique

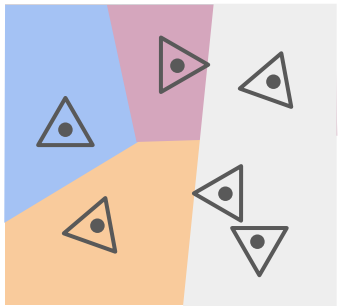

Abstain

Robust Models for Source Code

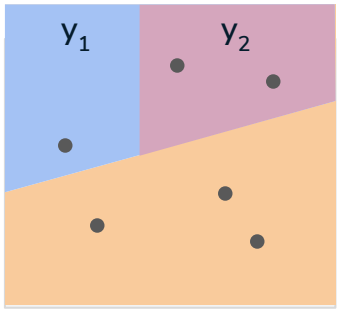

Standard Training

δ = hex ⟶ color

$\delta \in S(x)$

1st technique

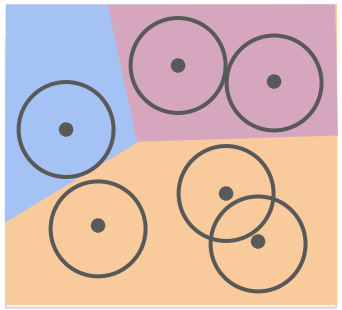

Adversarial Training

$\alpha(\mathbf{x} + \delta)$

2nd technique

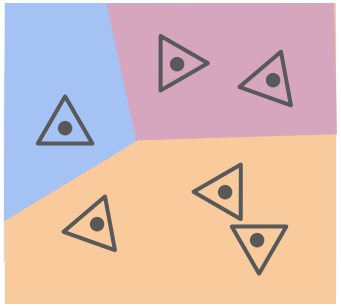

Representation Learning

3rd technique

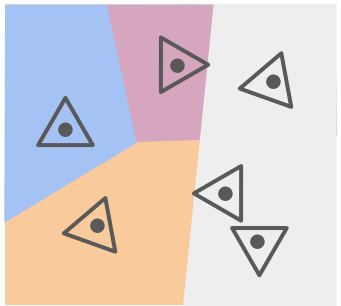

Abstain

Next technique?

- pre-processing
- compositionality
- normalization