

IvySyn

Automated Vulnerability Discovery in Deep Learning Frameworks

[Neophytos Christou](#)¹

[Di Jin](#)¹

[Vaggelis Atlidakis](#)¹

[Baishakhi Ray](#)²

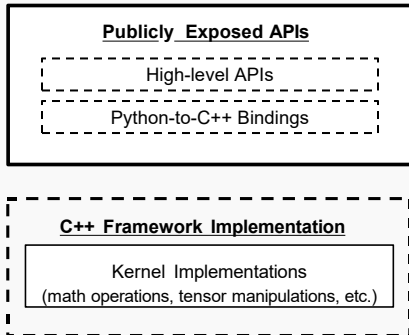
[Vasileios P. Kemerlis](#)¹

August 9, 2023

[¹Secure Systems Laboratory \(SSL\)](#)
[Department of Computer Science](#)
[Brown University](#)

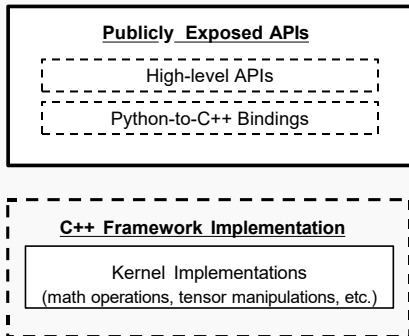
[²ARiSE Lab](#)
[Department of Computer Science](#)
[Columbia University](#)

DL Framework Architecture



- DL Frameworks offers several high-level APIs for performing DL-specific operations, usually written in managed languages, like Python.
- For performance, the essential parts of such operations(kernel) are implemented in low-level, memory-unsafe languages, like C/C++.

DL Framework Architecture



- Developer-Accessible
- High-level language (ex. Python)
- Memory-safe

DL Framework Architecture

Publicly Exposed APIs

High-level APIs

Python-to-C++ Bindings

- Memory-unsafe code region (raw memory operations)
- Not accessible directly
- BUT accessible and possibly exploitable through the Bindings and APIs

C++ Framework Implementation

Kernel Implementations
(math operations, tensor manipulations, etc.)

Motivation

- Core DL framework implementation → **memory-unsafe** languages
- In 2021–2022 alone, TensorFlow had more than **280 CVE numbers assigned for (potentially-exploitable) vulnerabilities** related to memory safety issues.

TFSA-2021-126	Use after free in boosted trees creation
TFSA-2021-125	Heap buffer overflow in <code>FractionalAvgPoolGrad</code>
TFSA-2021-124	Segfault and heap buffer overflow in <code>{Experimental,}DatasetToTFRecord</code>
TFSA-2021-123	Null pointer dereference in <code>UncompressElement</code>
TFSA-2021-122	Incorrect validation of <code>SaveV2</code> inputs
TFSA-2021-121	Null pointer dereference in <code>SparseTensorSliceDataset</code>
TFSA-2021-120	Bad alloc in <code>StringNGrams</code> caused by integer conversion

© IvySyn's Goals

© IvySyn's Goals

- ▶ **Automatically** uncover *memory safety* and *fatal runtime* errors in DL frameworks

© IvySyn's Goals

- ▶ **Automatically** uncover *memory safety* and *fatal runtime* errors in DL frameworks
- ▶ Help framework developers *identify* and *fix* the uncovered bugs

Goals and Past Approaches

© IvySyn's Goals

- ▶ **Automatically** uncover *memory safety* and *fatal runtime* errors in DL frameworks
- ▶ Help framework developers *identify* and *fix* the uncovered bugs

⚠ Past Approaches

Goals and Past Approaches

© IvySyn's Goals

- ▶ **Automatically** uncover *memory safety* and *fatal runtime* errors in DL frameworks
- ▶ Help framework developers *identify* and *fix* the uncovered bugs

⚠ Past Approaches

- ▶ Are **not** aimed at finding memory safety errors

Goals and Past Approaches

© IvySyn's Goals

- ▶ **Automatically** uncover *memory safety* and *fatal runtime* errors in DL frameworks
- ▶ Help framework developers *identify* and *fix* the uncovered bugs

⚠ Past Approaches

- ▶ Are **not** aimed at finding memory safety errors
- ▶ Are **not** *fully automated*
 - Custom fuzzing drivers, domain-expert annotations, ...

The general approach is based on these questions:

- Are there any offending inputs to low-level APIs that can trigger memory errors in native DL kernels?
- Are there any high-level APIs that can propagate the above offending inputs to low-level kernel code?

The general approach is based on these questions:

- Are there any offending inputs to low-level APIs that can trigger memory errors in native DL kernels?
- Are there any high-level APIs that can propagate the above offending inputs to low-level kernel code?



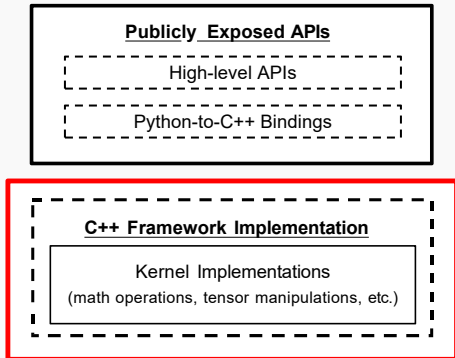
Bottom-up approach for testing DL framework implementations

< IvySyn's Approach

< IvySyn's Approach

- ▶ Fuzz the **native implementation** of DL frameworks

DL Framework Architecture



< IvySyn's Approach

- ▶ Fuzz the **native implementation** of DL frameworks
- ▶ **Automatically synthesize Proof-of-Vulnerability (PoV)** snippets

PoV synthesized by IvySyn

```
import tensorflow as tf

indices = tf.constant([], shape=[2,0],
                      dtype=tf.int64)
values = tf.constant([2,0,1], shape=[3],
                    dtype=tf.int64)
dense_shape = tf.constant([2,0,1], shape=[3],
                          dtype=tf.int64)
default_value = tf.constant(0, shape=[],
                            dtype=tf.int64)

tf.raw_ops.SparseFillEmptyRows(
    indices=indices,
    values=values,
    dense_shape=dense_shape,
    default_value=default_value)
```

PoV triggers crash!

```
$ python3 pov.py
segmentation fault (core dumped)
```

< IvySyn's Approach

- ▶ Fuzz the **native implementation** of DL frameworks
- ▶ **Automatically synthesize Proof-of-Vulnerability (PoV)** snippets

🏆 Achievements

📄 PoV synthesized by IvySyn

```
import tensorflow as tf

indices = tf.constant([], shape=[2,0],
                      dtype=tf.int64)
values = tf.constant([2,0,1], shape=[3],
                    dtype=tf.int64)
dense_shape = tf.constant([2,0,1], shape=[3],
                          dtype=tf.int64)
default_value = tf.constant(0, shape=[],
                           dtype=tf.int64)

tf.raw_ops.SparseFillEmptyRows(
    indices=indices,
    values=values,
    dense_shape=dense_shape,
    default_value=default_value)
```

🚩 PoV triggers crash!

```
$ python3 pov.py
segmentation fault (core dumped)
```

< IvySyn's Approach

- ▶ Fuzz the **native implementation** of DL frameworks
- ▶ **Automatically synthesize Proof-of-Vulnerability (PoV)** snippets

🏆 Achievements

- ▶ Uncovered **61** previously-unknown vulnerabilities

📄 PoV synthesized by IvySyn

```
import tensorflow as tf

indices = tf.constant([], shape=[2,0],
                      dtype=tf.int64)
values = tf.constant([2,0,1], shape=[3],
                    dtype=tf.int64)
dense_shape = tf.constant([2,0,1], shape=[3],
                          dtype=tf.int64)
default_value = tf.constant(0, shape=[],
                            dtype=tf.int64)

tf.raw_ops.SparseFillEmptyRows(
    indices=indices,
    values=values,
    dense_shape=dense_shape,
    default_value=default_value)
```

🚩 PoV triggers crash!

```
$ python3 pov.py
segmentation fault (core dumped)
```

< IvySyn's Approach

- ▶ Fuzz the **native implementation** of DL frameworks
- ▶ **Automatically synthesize Proof-of-Vulnerability (PoV)** snippets

🏆 Achievements

- ▶ Uncovered **61 previously-unknown vulnerabilities**
- ▶ Assigned with **39 unique CVEs**

📄 PoV synthesized by IvySyn

```
import tensorflow as tf

indices = tf.constant([], shape=[2,0],
                      dtype=tf.int64)
values = tf.constant([2,0,1], shape=[3],
                    dtype=tf.int64)
dense_shape = tf.constant([2,0,1], shape=[3],
                          dtype=tf.int64)
default_value = tf.constant(0, shape=[],
                            dtype=tf.int64)

tf.raw_ops.SparseFillEmptyRows(
    indices=indices,
    values=values,
    dense_shape=dense_shape,
    default_value=default_value)
```

🚩 PoV triggers crash!

```
$ python3 pov.py
segmentation fault (core dumped)
```

IvySyn Architecture -
Instrumentation

Evaluation

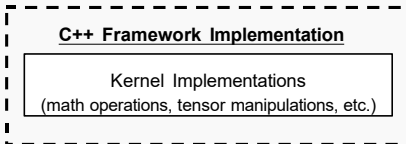
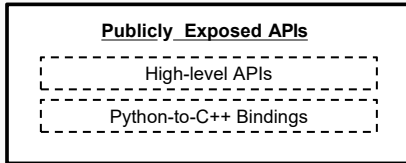
Conclusion

Future Plans

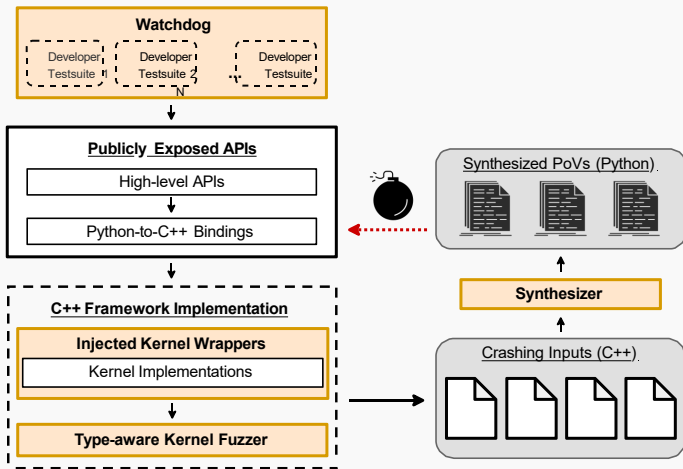
IvySyn:

- Focuses on two famous DL frameworks: TensorFlow and Pytorch.
- Assumes its target kernels are strongly-typed and avoid shared state.
- Assumes the existence of developer-provided unit tests.
- Leverages the DL framework mappings between high-level and low-level APIs.

DL Framework Architecture



IvySyn Architecture Overview



IvySyn Architecture → Extracting Kernel Implementations

In DL frameworks like PyTorch and TensorFlow every operation (Conv2D, ReLU, MatMul, etc.) must be registered with the framework.

Example (Tensor-flow style):

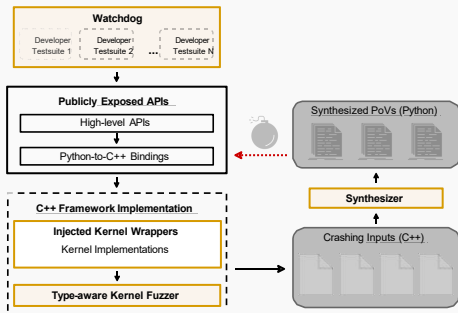
```
REGISTER_KERNEL_BUILDER(  
    Name("Conv2D").Device(DEVICE_CPU),  
    Conv2D0p);
```

This gives IvySyn structured metadata for fuzzing:

```
Kernel: Conv2D  
Arguments:  
  Tensor input  
  Tensor weight  
  int stride  
  int padding  
Device: CPU
```

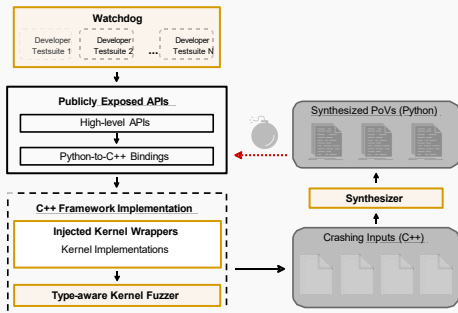
This is enough to construct type-aware mutations and start the fuzzing sessions.

- ▶ IvySyn **automatically** wraps each framework's *kernels* with fuzzing drivers

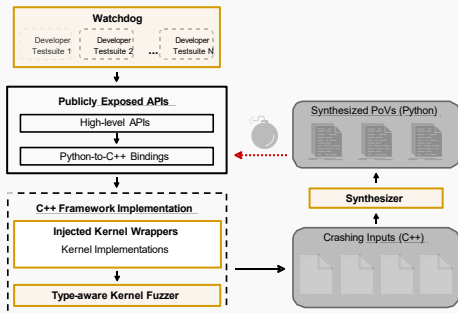


- ▶ IvySyn **automatically** wraps each framework's *kernels* with fuzzing drivers

How?

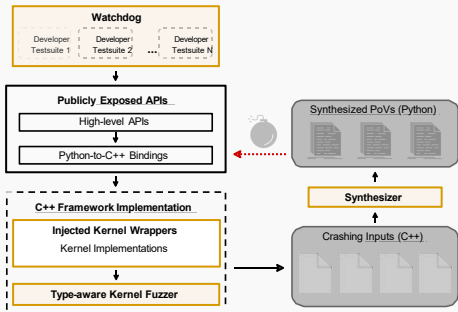


- ▶ IvySyn **automatically** wraps each framework's *kernels* with fuzzing drivers
- ▶ Original kernel `foo()` ...→ `do_foo()`
- ▶ Creates instrumented kernel named `foo()`

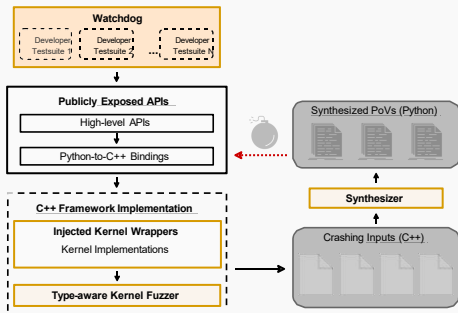


IvySyn Architecture → Kernel Instrumentation

```
1 Tensor cosine_similarity(const Tensor& x1,
2   const Tensor& x2, int64_t dim, double eps){
3
4   if (fuzzing::already_fuzzing ||
5       fuzzing::was_fuzzed("cosine_similarity"))
6       // Return the original function result
7       return do_cosine_similarity(x1, x2,
8         dim, eps);
9   fuzzing::already_fuzzing = true;
10
11   Tensor retval, x1_fuzz, x2_fuzz;
12   int64_t dim_fuzz; double eps_fuzz;
13   std::vector<std::string> types =
14     {"Tensor&", "Tensor&",
15      "int64_t", "double"};
16   std::vector<void*> args{};
17   args.push_back((void*) &x1);
18   ... // Repeat for x2, dim, and eps
19   // Initialize fuzzer with the original
20   // args and types
21   fuzzing::Fuzzer fuzzer =
22     fuzzing::Fuzzer("cosine_similarity",
23       types, args);
24
25   while (fuzzer.has_more_mutations(true)){
26     // Get the next combination of inputs
27     x1_fuzz = fuzzer.get_next_mut_tensor();
28     ... // Repeat for x2, dim, and eps
29     try {
30       // Invoke the original function
31       fuzzer.mut_start_time();
32       do_cosine_similarity(x1_fuzz,
33         x2_fuzz, dim_fuzz, eps_fuzz);
34       fuzzer.mut_end_time(false);
35     } catch (...) { ... }
36   }
37   fuzzing::already_fuzzing = false;
38
39   // Return the original function result
40   return do_cosine_similarity(x1, x2,
41     dim, eps);
42 }
```

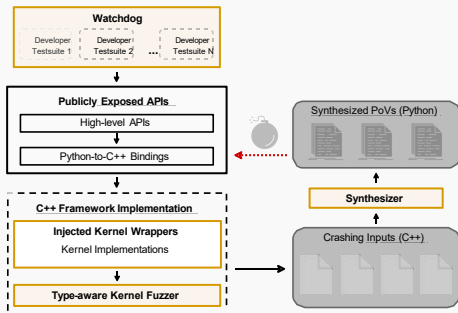


- ▶ IvySyn runs DL frameworks' **developer-provided unit tests**



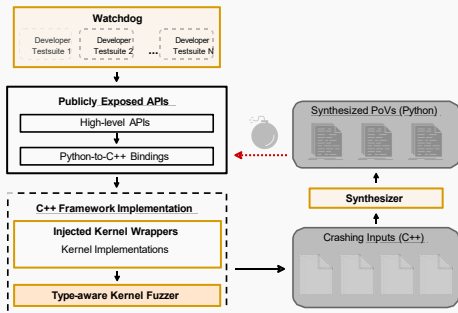
IvySyn Architecture → Force-executing Instrumented Kernels

- ▶ IvySyn runs DL frameworks' **developer-provided unit tests**
- ▶ Execution reaches the *instrumented kernels* ...



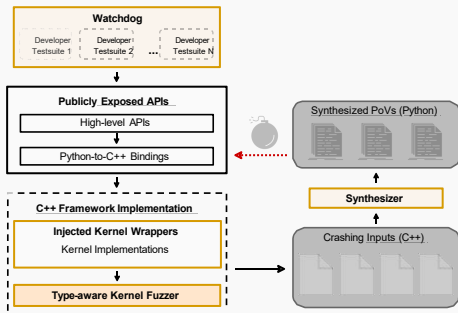
IvySyn Architecture → Force-executing Instrumented Kernels

- ▶ IvySyn runs DL frameworks' **developer-provided unit tests**
- ▶ Execution reaches the *instrumented* kernels ...
- ▶ ...and bootstraps a **fuzzing session**

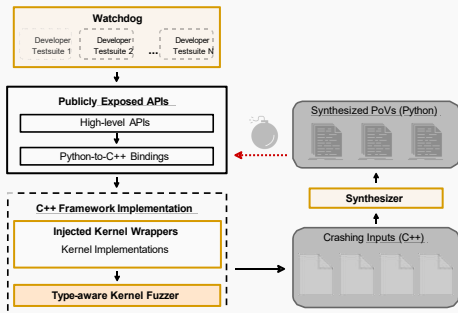


- ▶ IvySyn runs DL frameworks' **developer-provided unit tests**
- ▶ Execution reaches the *instrumented* kernels ...
- ▶ ...and bootstraps a **fuzzing session**

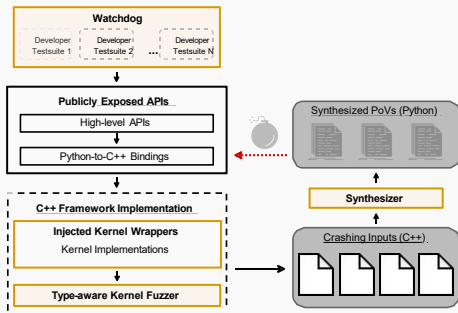
Why does IvySyn need the unit tests?



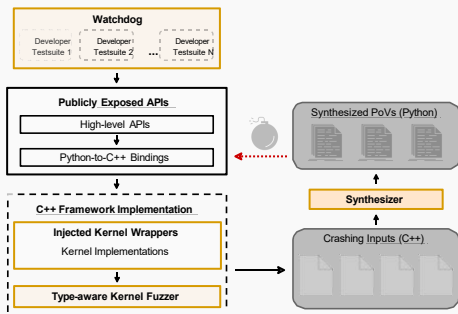
- ▶ Performs **type-aware** mutations based on the original argument types



- ▶ Performs **type-aware** mutations based on the original argument types
- ▶ Logs *native crashing inputs* in **crash-reports**

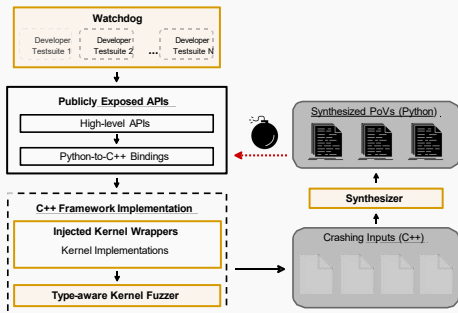


- ▶ Logged crash-reports are fed into IvySyn's **synthesizer**

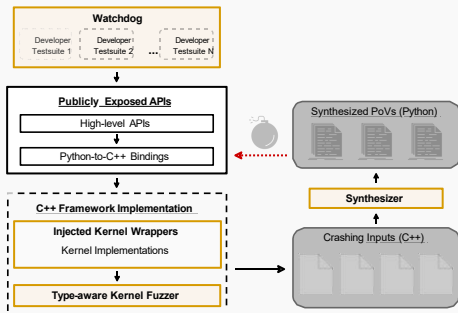


IvySyn Architecture → PoV Synthesis

- ▶ Logged crash-reports are fed into IvySyn's **synthesizer**
- ▶ The synthesizer generates *Proof-of-Vulnerability (PoV)* snippets



- ▶ Logged crash-reports are fed into IvySyn's **synthesizer**
- ▶ The synthesizer generates *Proof-of-Vulnerability (PoV)* snippets
- ▶ The PoVs trigger the native crashes from **publicly exposed Python APIs**



Crash-report produced by IvySyn

```
# SparseFillEmptyRowsOp  
Tensor<type: int64 shape: [2,0] values: >  
Tensor<type: int64 shape: [3] values: 2 0 1>  
Tensor<type: int64 shape: [3] values: 2 0 1>  
Tensor<type: int64 shape: [] values: 0>
```

Crash-report produced by IvySyn

```
# SparseFillEmptyRowsOp
Tensor<type: int64 shape: [2,0] values: >
Tensor<type: int64 shape: [3] values: 2 0 1>
Tensor<type: int64 shape: [3] values: 2 0 1>
Tensor<type: int64 shape: [] values: 0>
```

Corresponding PoV synthesized by IvySyn

```
import tensorflow as tf

indices = tf.constant([], shape=[2,0],
                      dtype=tf.int64)
values = tf.constant([2,0,1], shape=[3],
                    dtype=tf.int64)
dense_shape = tf.constant([2,0,1], shape=[3],
                          dtype=tf.int64)
default_value = tf.constant(0, shape=[],
                            dtype=tf.int64)

tf.raw_ops.SparseFillEmptyRows(
    indices=indices,
    values=values,
    dense_shape=dense_shape,
    default_value=default_value)
```

Crash-report produced by IvySyn

```
# SparseFillEmptyRowsOp
Tensor<type: int64 shape: [2,0] values: >
Tensor<type: int64 shape: [3] values: 2 0 1>
Tensor<type: int64 shape: [3] values: 2 0 1>
Tensor<type: int64 shape: [] values: 0>
```

PoV triggers crash!

```
$ python3 pov.py
segmentation fault (core dumped)
```

Corresponding PoV synthesized by IvySyn

```
import tensorflow as tf

indices = tf.constant([], shape=[2,0],
                      dtype=tf.int64)
values = tf.constant([2,0,1], shape=[3],
                    dtype=tf.int64)
dense_shape = tf.constant([2,0,1], shape=[3],
                          dtype=tf.int64)
default_value = tf.constant(0, shape=[],
                            dtype=tf.int64)

tf.raw_ops.SparseFillEmptyRows(
    indices=indices,
    values=values,
    dense_shape=dense_shape,
    default_value=default_value)
```

IvySyn Architecture -
Instrumentation

Evaluation

Conclusion

Future Plans





TensorFlow



PyTorch

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?



- Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?
- Q2.** Is *IvySyn* effective in leveraging crashing inputs to synthesize PoVs?



TensorFlow



PyTorch

- Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?
- Q2. Is *IvySyn* effective in leveraging crashing inputs to synthesize PoVs?
- Q3.** Which *IvySyn* mutations are the most successful in uncovering memory errors?



TensorFlow



PyTorch

- Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?
- Q2. Is *IvySyn* effective in leveraging crashing inputs to synthesize PoVs?
- Q3. Which *IvySyn* mutations are the most successful in uncovering memory errors?
- Q4.** What are the security ramifications of the PoVs synthesized with *IvySyn*?

Evaluation → IvySyn vs Atheris

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?

- ▶ Compared *IvySyn's efficiency* at uncovering crashes against *Atheris*[†]

[†]*Atheris: A Coverage-Guided, Native Python Fuzzer.* Google.

Evaluation → IvySyn vs Atheris

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?

- ▶ Compared *IvySyn*'s *efficiency* at uncovering crashes against *Atheris*[†]
- ▶ Leveraged *IvySyn*'s **argument logging functionality** and **synthesizer** to generate *fuzzing drivers* for *Atheris*

[†]*Atheris: A Coverage-Guided, Native Python Fuzzer*. Google.

Evaluation → IvySyn vs Atheris

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?

- ▶ Compared *IvySyn*'s *efficiency* at uncovering crashes against *Atheris*[†]
- ▶ Leveraged *IvySyn*'s **argument logging functionality** and **synthesizer** to generate *fuzzing drivers* for *Atheris*
- ▶ Generated two different variants of drivers

[†]*Atheris: A Coverage-Guided, Native Python Fuzzer*. Google.

Evaluation → IvySyn vs Atheris

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?

- ▶ Compared *IvySyn*'s *efficiency* at uncovering crashes against *Atheris*[†]
- ▶ Leveraged *IvySyn*'s **argument logging functionality** and **synthesizer** to generate *fuzzing drivers* for *Atheris*
- ▶ Generated two different variants of drivers

Atheris[†]

- Drivers **without** type awareness
- *Atheris* randomly chooses argument types

[†]*Atheris: A Coverage-Guided, Native Python Fuzzer*. Google.

Evaluation → IvySyn vs Atheris

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?

- ▶ Compared *IvySyn*'s *efficiency* at uncovering crashes against *Atheris*[†]
- ▶ Leveraged *IvySyn*'s **argument logging functionality** and **synthesizer** to generate *fuzzing drivers* for *Atheris*
- ▶ Generated two different variants of drivers

Atheris⁺

- Drivers **without** type awareness
- *Atheris* randomly chooses argument types



Atheris⁺⁺

- Drivers **with** type awareness
- The drivers provide *Atheris* with the proper argument types

[†]*Atheris: A Coverage-Guided, Native Python Fuzzer.* Google.


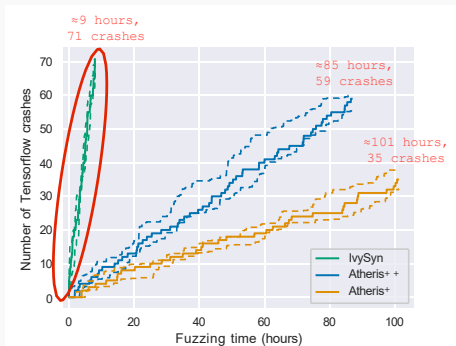
Evaluation → IvySyn vs Atheris (cont'd)

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?



Evaluation → IvySyn vs Atheris (cont'd)

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?



► *IvySyn* uncovers **more crashes** than *Atheris*, and does so **faster**

Q1. Is *IvySyn* efficient in uncovering crashing inputs over time?

Number of crashes found by IvySyn vs Atheris (over 5 iterations)

	Fuzzer	TensorFlow	PyTorch
Total Crashes	Atheris ⁺	47	9
	Atheris ⁺⁺	64	18
	IvySyn	80	25
Union	All	87	30

- ▶ IvySyn uncovers **more crashes** than Atheris, and does so **faster**

Q2. Is *IvySyn* effective in leveraging crashing inputs to synthesize PoVs?

- ▶ Compared *IvySyn*'s *effectiveness* at synthesizing PoVs against the *semi-automated* DocTer[‡] tool

[‡]*DocTer: Documentation-Guided Fuzzing for Testing Deep Learning API Functions*
Xie et al.

Q2. Is *IvySyn* effective in leveraging crashing inputs to synthesize PoVs?

- ▶ Compared *IvySyn*'s *effectiveness* at synthesizing PoVs against the *semi-automated* DocTer[‡] tool

Synthesized PoVs (*IvySyn* vs DocTer) (over 10 iterations)

Fuzzer	Total		Median		Median Running Time (mins)	
	TensorFlow	PyTorch	TensorFlow	PyTorch	TensorFlow	PyTorch
DocTer	16	9	12	7	199	736
IvySyn	19	14	15	11	184	569

- ▶ *IvySyn* synthesizes **more PoVs** than DocTer, *without* manual effort

[‡]*DocTer: Documentation-Guided Fuzzing for Testing Deep Learning API Functions.*
Xie et al.

Q2. Is *IvySyn* effective in leveraging crashing inputs to synthesize PoVs?

Accumulated *IvySyn* results

Framework	Fuzzed Kernels	Unique Crashes	Synthesized PoVs
TensorFlow	412	103	86 / 103 (83%)
PyTorch	747	81	49 / 81 (60%)
All	1159	184	135 / 184 (73%)

- ▶ *IvySyn* synthesized 135 PoVs and was attributed with 39 CVEs

Evaluation → Effectiveness per Mutation Type

Q3. Which *IvySyn* mutations are the most successful in uncovering memory errors?

Number of PoVs per mutation type

IvySyn Type-aware Mutation Type	Total PoVs
Tensors with random dimension sizes	46
Tensors with extreme values	25
Permutations of original arguments	19
Zero values	12
Lists with extreme values	13
Tensors with empty shape	8
Extreme values in primitive types	6
Empty lists	3
Deep tensors	3

Evaluation → Effectiveness per Mutation Type

Q3. Which *IvySyn* mutations are the most successful in uncovering memory errors?

Number of PoVs per mutation type

IvySyn Type-aware Mutation Type	Total PoVs
Tensors with random dimension sizes	46
Tensors with extreme values	25
Permutations of original arguments	19
Zero values	12
Lists with extreme values	13
Tensors with empty shape	8
Extreme values in primitive types	6
Empty lists	3
Deep tensors	3

- **DL-specific** (e.g., tensor) mutations are especially effective

Q4. What are the security ramifications of the PoVs synthesized with *IvySyn*?

- *IvySyn* successfully synthesized PoVs for 135 of the total 184 TensorFlow and PyTorch crashes found.

Q4. What are the security ramifications of the PoVs synthesized with *IvySyn*?

- *IvySyn* successfully synthesized PoVs for 135 of the total 184 TensorFlow and PyTorch crashes found.
- To what extent these PoVs could be abused to corrupt memory locations, leak program contents, or cause the respective runtime environment to crash?

Q4. What are the security ramifications of the PoVs synthesized with *IvySyn*?

- *IvySyn* successfully synthesized PoVs for 135 of the total 184 TensorFlow and PyTorch crashes found.
- To what extent these PoVs could be abused to corrupt memory locations, leak program contents, or cause the respective runtime environment to crash?

Framework	SIGABRT	SIGFPE	SIGSEGV
TensorFlow	56	4	26
PyTorch	-	16	33

Q4. What are the security ramifications of the PoVs synthesized with *IvySyn*?

- *IvySyn* successfully synthesized PoVs for 135 of the total 184 TensorFlow and PyTorch crashes found.
- To what extent these PoVs could be abused to corrupt memory locations, leak program contents, or cause the respective runtime environment to crash?

Framework	SIGABRT	SIGFPE	SIGSEGV
TensorFlow	56	4	26
PyTorch	-	16	33

- Most crashes do not cause memory corruption(floating-point exceptions, aborts), but there are still security-relevant threats(possible DoS).

Q4. What are the security ramifications of the PoVs synthesized with *IvySyn*?

- *IvySyn* successfully synthesized PoVs for 135 of the total 184 TensorFlow and PyTorch crashes found.
- To what extent these PoVs could be abused to corrupt memory locations, leak program contents, or cause the respective runtime environment to crash?

Framework	SIGABRT	SIGFPE	SIGSEGV
TensorFlow	56	4	26
PyTorch	-	16	33

- Most crashes do not cause memory corruption(floating-point exceptions, aborts), but there are still security-relevant threats(possible DoS).
- Majority of segfaults(memory read and write operations at controlled addresses) are security critical:
By controlling the addresses of these operations, an attacker could leverage them to gain arbitrary read/write primitives.

IvySyn Architecture -
Instrumentation

Evaluation

Conclusion

Future Plans

- ▶ Fully-automated framework
 - Perform type-aware, DL-specific mutations

- ▶ Fully-automated framework
 - Perform type-aware, DL-specific mutations
 - Fuzz *native implementation* of DL frameworks

- ▶ **Fully-automated** framework
 - Perform **type-aware, DL-specific** mutations
 - **Fuzz** *native implementation* of DL frameworks
 - **Synthesize** PoVs to trigger detected crashes from Python

- ▶ **Fully-automated** framework
 - Perform **type-aware, DL-specific** mutations
 - **Fuzz native implementation** of DL frameworks
 - **Synthesize** PoVs to trigger detected crashes from Python
- ▶ Identified **61 previously-unknown vulnerabilities**

- ▶ Fully-automated framework
 - Perform type-aware, DL-specific mutations
 - Fuzz native implementation of DL frameworks
 - Synthesize PoVs to trigger detected crashes from Python
- ▶ Identified 61 previously-unknown vulnerabilities
- ▶ Assigned with 39 unique CVEs

IvySyn Architecture -
Instrumentation

Evaluation

Conclusion

Future Plans

Scope Expansion

- Extend to additional DL frameworks and hardware backends.
- Explore applicability to non-DL codebases.

Scope Expansion

- Extend to additional DL frameworks and hardware backends.
- Explore applicability to non-DL codebases.

Technical Advancements

- Support stateful kernels or kernels with non-{tensor, array, scalar} arguments types.
- Coverage-guided, type-aware kernel fuzzing (combine IvySyn's kernel hooks with Atheris-style feedback).
- GPU-Aware and Race-Condition Fuzzing.

THANK YOU