

# A Static Analyzer for Detecting Tensor Shape Errors in Deep Neural Network Training Code

Ho Young Jhoo  
Seoul National University  
Seoul, South Korea  
hoyoung.jhoo@sf.snu.ac.kr

Sehoon Kim  
Seoul National University  
Seoul, South Korea  
shkim@ropas.snu.ac.kr

Woosung Song  
Seoul National University  
Seoul, South Korea  
lego0901@gmail.com

Kyuyeon Park  
Seoul National University  
Seoul, South Korea  
kypark@ropas.snu.ac.kr

DongKwon Lee  
Seoul National University  
Seoul, South Korea  
dklee@ropas.snu.ac.kr

Kwangkeun Yi  
Seoul National University  
Seoul, South Korea  
kwang@ropas.snu.ac.kr

## ABSTRACT

We present an automatic static analyzer PyTea that detects tensor-shape errors in PyTorch code. The tensor-shape error is critical in the deep neural net code; much of the training cost and intermediate results are to be lost once a tensor shape mismatch occurs in the midst of the training phase. Given the input PyTorch source, PyTea statically traces every possible execution path, collects tensor shape constraints required by the tensor operation sequence of the path, and decides if the constraints are unsatisfiable (hence a shape error can occur). PyTea’s scalability and precision hinges on the characteristics of real-world PyTorch applications: the number of execution paths after PyTea’s conservative pruning rarely explodes and loops are simple enough to be circumscribed by our symbolic abstraction. We tested PyTea against the projects in the official PyTorch repository and some tensor-error code questioned in the StackOverflow. PyTea successfully detects tensor shape errors in these codes, each within a few seconds.

## CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging**.

## KEYWORDS

static analysis, error detection, tensor shape mismatch, neural networks, SMT solver, Python, PyTorch

## 1 INTRODUCTION

### 1.1 Our Goal

*Tensor shape mismatch* is a critical bug in deep neural network machine learning applications. Training a neural network is an expensive process that intends to terminate only when it finishes processing a huge amount of data through a sequence of tensor operations. In the middle of this time-consuming training process, if the shape of an input datum failed to fit with a tensor operation, the whole process abruptly stops wasting the entire training cost spent thus far, losing the trained, if any, intermediate result.

Our goal is to automatically predict at compile-time such run-time tensor-shape mismatch errors in PyTorch neural network training code.

### 1.2 Structure of PyTorch Programs

Contemporary machine learning frameworks such as PyTorch [23], TensorFlow [8], and Keras [13] use Python APIs to build neural networks. Training a neural network with such frameworks is mostly patterned after a standard procedure which is illustrated in Figure 1. Typical PyTorch neural network training code can be divided into four stages. Figure 2 shows a code example, a simplified image classification code taken from the official PyTorch MNIST classification example [5]. We first define the series of neural network layers and make them into a single neural network module. To correctly assemble the layers, the returned tensor of the former layer must satisfy the input requirements of the next layer. We will see those requirements from the next section. The network is instantiated with some initialization parameters called hyperparameter, e.g., the number of hidden layers. Next, the input dataset is preprocessed and adjusted to the requirements of the network. Every dataset is cut into smaller same-sized chunks (called minibatches) from this stage. Finally, the main loop starts, and the minibatches are sequentially fed to the network. One epoch means a single loop that an entire dataset is passed to the network, and the number of epochs (datasets) usually differs depending on the purpose and structure of the neural network. Including the number of epochs, the numbers of iterations in the training code are determined to be constants in most cases, except the main training loop which depends on the size of a dataset.

### 1.3 Tensor Shape Errors

Figure 3 presents the typical type of tensor shape errors, which are slight modifications of Figure 2. From the first example, the second `Linear` layer (line 8), which multiplies the input with  $80 \times 10$ -matrix, requires a specific shape of a tensor as an input. The first layer (line 6), however, returns a wrong-shaped tensor, and the overall pipeline will malfunction. This kind of error is called *tensor shape mismatch* error, simply, shape error.

Shape error is rather hard to manually find, only to be detected by running the program with an actual input. Indeed, the most notorious error for machine learning engineers is the error that can only be occurred after an immense amount of machine-hours.

Figure 3(b) shows another example. Its declaration of training data loader (line 14) hides a shape error. `DataLoader` class slices the dataset sequentially by `batch_size` and passes it to the model. If the

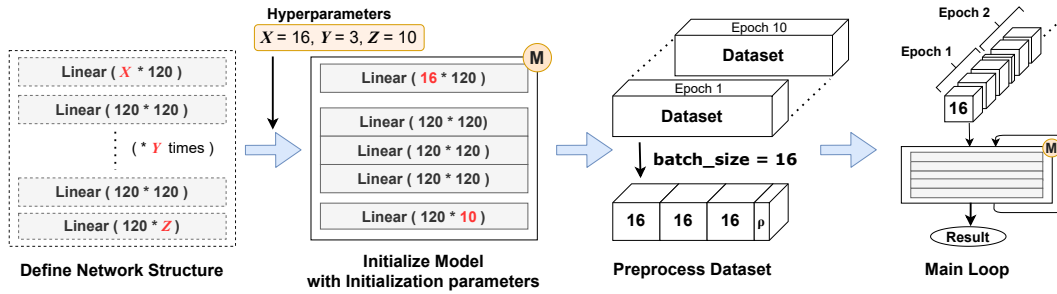


Figure 1: Typical structure of neural network training code in PyTorch.

```

1  ## 1. DEFINE NETWORK STRUCTURE
2  class Net(nn.Module):
3      def __init__(self, out_classes):
4          super(Net, self).__init__()
5          self.layers = nn.Sequential(
6              nn.Linear(28 * 28, 120),
7              nn.ReLU(),
8              nn.Linear(120, out_classes)
9          )
10
11     def forward(self, x):
12         x = x.reshape(x.shape[0], -1)
13         x = self.layers(x)
14         return x
15
16     ## 2. INITIALIZE MODEL
17     model = Net(out_classes=10)
18
19     ## 3. PREPROCESS DATASET
20     data = dataset.MNIST('./data', train=True,
21                         transform=[ToTensor()])
22     loader = DataLoader(data, batch_size=16)
23
24     ## 4. RUN MAIN LOOP
25     for epoch in range(10):
26         for batch, label in loader:
27             # model(batch) == model.forward(batch)
28             output = model(batch)
29             loss = F.nll_loss(output, label)
30             loss.backward()

```

Figure 2: Basic PyTorch training code.

total length of the dataset is not divisible by batch\_size, however, the size of the residual minibatch will be the non-zero remainder of the total length. See line 16: because the third parameter drop\_last is missing, the model assumes a consistent batch size (lines 10 and 6) hence the program will crash from the residual minibatch, losing the whole training hours. The recent massive networks like GPT-3 [11] require more than hundreds of machine-hours to train. This type of error must be noticed before its run.

Figure 3(c) illustrates another shape error that can be arisen from a dataset, not a structure of the model. It does not take input from the pre-defined MNIST dataset but reads an image from a file. If the read image is RGB, which has 3xHxW dimensions, it will not fit into the reshape method that requires a tensor of 28x28-elements. That means we have to convert it to a monochrome image before

```

1  class Net(nn.Module):
2      def __init__(self):
3          super(Net, self).__init__()
4          self.layers = nn.Sequential(
5              ## 'B' represents batch size
6              ## [B x 784] * [784 x 120] -> [B x 120]
7              nn.Linear(28 * 28, 120),
8              ## [B x 120] -> [B x 120]
9              nn.ReLU(),
10             ## [B x 120] * [80 x 10] -> ERROR!
11             nn.Linear(80, 10))

```

(a) Error on the network structure.

```

1  class Net(nn.Module):
2      def __init__(self, batch_size):
3          self.batch_size = batch_size
4          # ...
5      def forward(self, x):
6          x = x.reshape(self.batch_size, -1)
7          # ...
8
9      ## some models may require exact batch size
10     model = Net(batch_size=64)
11
12     ## POTENTIAL_ERROR 1:
13     ## argument 'drop_last=True' is essential
14     loader = DataLoader(data, batch_size=64)
15     # loader = DataLoader(data, batch_size=64,
16     #                       drop_last=True)
17
18     for epoch in range(10):
19         for batch, label in loader:
20             out = model(batch)
21             ## ERROR ON THE LAST MINIBATCH
22             ## last batch size: 32 (!= 64)

```

(b) Error on the last minibatch.

```

1  ## POTENTIAL_ERROR 2: channel size can be 3
2  img = PIL.Image.open('./image.png').resize([28, 28])
3  # img = img.convert('L')
4
5  ## ERROR WHEN THE IMAGE IS RGB.
6  tensor = to_tensor(img).reshape(28 * 28)
7  out = model(tensor)

```

(c) Insufficient data preprocessing.

Figure 3: Various type of tensor shape errors.

feeding it to the network. Even though it had been successfully tested with monochrome images, there can be a user who tests it with an RGB image, crashing the execution of the code.




Figure 4: Overall architecture of PyTea.

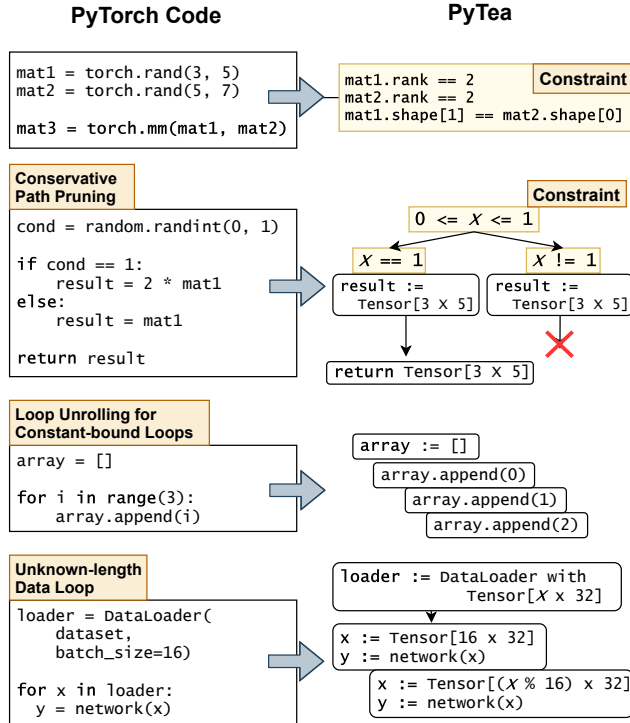


Figure 5: Constraint generation example.

Though several works [14, 16, 20, 22, 26] have reported tools to detect the shape mismatch errors of machine learning libraries, especially for TensorFlow [8], none of them have presented any static analysis tool that statically detects the shape errors for realistic Python ML applications. Real-world machine learning applications heavily utilize third-party libraries, external datasets, and configuration parameters, and handle their controls with subtle branch conditions and loops, but the existing tools still lack in supporting some of these elements and thus they fail to analyze even a simple ML application. To ensure that the shape error will not happen for *any* input data, we should statically infer a precise yet conservative range of each tensor shape and track its transformations through all possible execution paths.

```

1 class RandBlock(nn.Module):
2     def __init__(self):
3         super(RandBlock, self).__init__()
4         self.layer = nn.Linear(32, 32)
5
6     def forward(self, x):
7         rand_num = random.randint(0, 1)
8
9         if rand_num == 1:
10            result = self.layer(x)
11        else:
12            result = x
13
14        return result
15
16 model = nn.Sequential(
17     [RandBlock() for _ in range(24)])
    
```

Figure 6: Path explosion example.

## 2 OVERVIEW OF PYTEA ANALYZER

To find out shape errors before runtime, we present a static analyzer **PyTea** (PyTorch Tensor Error Analyzer). PyTea statically scans PyTorch applications and detects possible shape errors. PyTea analyzes full training and evaluation paths of the real-world Python/PyTorch applications with additional data processing and mixed usage of other libraries (e.g., Torchvision [4], NumPy [19], PIL [1])

Figure 4 illustrates the overall architecture of PyTea analyzer. It first translates the original Python codes into a kernel language, PyTea Internal Representation (PyTea IR). Then, it tracks every possible execution path of the translated IR and collects the constraints regarding tensor shapes that dictate the conditions for the code to run without a shape error. The collected constraint sets are given to Satisfiability Modulo Theories (SMT) solver Z3 [15] to judge that those constraints are satisfiable for every possible input shape. Following the result of the solver, PyTea concludes which path contains a shape error or not. If the constraint-solving by Z3 takes too much time, PyTea stops and tells "don't know".

### 2.1 Assumptions

Given the typical structure of PyTorch neural network training code (Section 1.2), we assume for the PyTea's input the followings about the PyTorch deep neural network training code:

- A1** Other than the training or evaluation dataset, every input value required to execute the code is injected by command-line arguments.

- A2 There is no infinite loop and recursion. We assume that every loop bound except for the datasets will be fixed to a constant.
- A3 The unknown loop bound for the datasets is only for the size of each dataset in an epoch, and every iteration is either with a fixed-sized minibatch of the dataset or with a smaller, residual minibatch.
- A4 We assume that string-manipulation expressions have no effect on tensor shapes.

These assumptions are based on our observations that most PyTorch networks and codes can be statically determined to fixed structures once we give precise command-line arguments. Real-world PyTorch applications mostly construct their structures by command-line arguments or external configuration files like JSON files. Therefore, PyTea chooses to analyze programs only with exact command-line arguments.

For a few networks that are not resolved to a single fixed structure, we consider all possible structures. The number of the possible structures is to be controlled by our path-pruning technique, and sometimes, for an inevitable case, by timeout.

## 2.2 Handling path explosions

The number of possible paths is exponential to the number of branches in sequence. For some complex neural networks, such path explosion is possible. For example, Neural Architecture Search [28] or Networks with Stochastic Depth [21] have branches inside the network themselves. Figure 6 shows a representative path explosion case that utilizes a runtime random variable. We can notice that the feed-forward function (`forward(self, x)`) has two execution paths in its body. The final structure of the network is made with 24 same blocks (line 17), which makes 16M paths.

We handle this exponential cost blow-up by means of conservative path-pruning and simple-minded timeouts. If we can find that the result of the binding scope of that feed-forward function is pure (i.e., do not change any global value), and its bounded value is indeed equal for every path and not related with the branch conditions, we then safely ignore other paths except for one. If a path explosion arises even if using this method, we then use a timeout. See Section 3.2.3 for more details.

## 2.3 Handling Loops

For the loops in typical PyTorch neural network programs, as we discussed in Section 1.2 and accordingly assumed in Section 2.1, we do not need the full power of static analysis [25]. PyTea unrolls constant-bound loops (Assumption A2 in Section 2.1) and analyzes their straight-line code version.

For the unknown-bound loops for datasets, PyTea analyzes the loop body for just two cases with the aforementioned assumption A3. One is for the loop with a fixed-sized regular minibatch of an epoch. The other is for the loop with the residual minibatch. For example, see code in Figure 5. For the third code box of Figure 5, we can unroll the loop expression to 3 same expressions. If we do not know the length of the dataset, such as the fourth code box of Figure 5, we use assumption A3 and consider only two cases for the two different sizes of minibatches.

### Expression

$E$	$\rightarrow$	$n \in \mathbb{Z} \mid T \mid F \mid x$ (variable)
		$\mid \text{let } xEE$
		$\mid \text{if } EEE$
		$\mid E \text{ bop } E$
		$\mid \text{tensor-expr}$
$\text{bop}$	$\rightarrow$	$\text{numeric-op} \mid \text{compare-op}$
$\text{numeric-op}$	$\rightarrow$	$+ \mid - \mid * \mid \dots$
$\text{compare-op}$	$\rightarrow$	$< \mid = \mid \dots$
$\text{tensor-expr}$	$\rightarrow$	$\text{mm } EE$
		$\mid \text{reshape } EEE$
		$\mid \text{readImage} \mid \dots$

Figure 7: Abstract syntax of PyTea IR.<sup>1</sup>

## 3 ANALYSIS STEPS

### 3.1 PyTea IR




Figure 8: A tensor that has shape (2, 3, 4). The rank of this tensor is 3, and each dimension has size 2, 3, and 4.

As the first step of the analysis, the input Python code is translated into the kernel language, *PyTea IR*. See Figure 7. PyTorch APIs are translated into tensor expressions that only define shape transformations, which PyTea IR focuses on.

The second step of the analysis is to scan the PyTea IR code and generate constraints.

### 3.2 Constraint generation

*Constraints* are the conditions required by a PyTorch application so that it can be executed without any tensor shape error. For example, two operands of a matrix multiplication operation must share the same dimension. For each tensor operation (`mm`, `reshape`, `readImage`, etc. of Figure 7), the shape of the input tensor must obey the requirement of the corresponding operation.

Figure 9 shows the abstract syntax of the constraints. *Value expression* represents the value of PyTea IR expressions, which can

<sup>1</sup>For the explanatory purpose, we did not include function calls and definitions. See supplementary material for detailed definitions of PyTea IR. Currently, we implemented 34 basic tensor expressions, and every other PyTorch API has been constructed with the basic expressions. The basic expressions are as following: `Torch.__init__`, `Torch.__getitem__`, `isSameShape`, `scalar`, `identity`, `broadcast`, `matmul`, `mm`, `bmm`, `item`, `repeat`, `expand`, `expand_as`, `transpose`, `reduce`, `topk`, `view`, `conv2d`, `conv_transpose2d`, `pool2d`, `batchnorm2d`, `cross_entropy`, `cat`, `stack`, `unsqueeze`, `squeeze`, `diag`, `flatten`, `narrow`, `pixel_shuffle`, `layer_norm`, `pad`, `adaptive`, `interpolate`.

<i>Constraint</i>	
$c \rightarrow$	$c \wedge c$
	$c \vee c$
	$\neg c$
	$e_b$
	$e = e$
	$e_n < e_n$
	$\forall \alpha_n \in [e_n, e_n].c$ (c is true forall integer $\alpha_n$ in the interval)
<i>Value Expr</i>	
$e \rightarrow$	$e_s   e_n   e_b$ (shape, number, or boolean)
<i>Shape Expr</i>	
$e_s \rightarrow$	$(e_n, \dots, e_n)$ (tensor shape)
	$\alpha_s$ (unknown shape)
	$e_s[e_n : e_n]$ (shape slicing)
	$e_s @ e_s$ (shape concat)
<i>Number Expr</i>	
$e_n \rightarrow$	$n$ (const number)
	$\alpha_n$ (unknown number)
	$e_n \text{ bop } e_n$ (binary operator)
	$\text{rank}(e_s)$ (rank of shape)
	$e_s[e_n]$ ( $e_n$ -th dimension of shape $e_s$ )
	$\prod e_s$ (number of elements in tensor of shape $e_s$ )
$\text{bop} \rightarrow$	$+   -   *   \dots$
<i>Boolean Expr</i>	
$e_b \rightarrow$	$\text{True}   \text{False}$
	$\alpha_b$ (unknown boolean)
	$e_b \wedge e_b$ (conjunction)
	$e_b \vee e_b$ (disjunction)
	$\neg e_b$ (negation)
	$e = e$ (equality)
	$e_n < e_n$ (less than)

Figure 9: Abstract syntax of constraints.

be used inside shape constraints. When PyTea analyzes a PyTea IR, it traces tensor shapes and primitive values of Python and constructs symbolic value expressions. *Shape expression* represents the shape of tensors, which is basically a tuple of integers  $(e_n, \dots, e_n)$ . Figure 8 shows an example of a tensor with a shape  $(2, 3, 4)$ . Each integer is a dimension size. We call the number of dimensions as a *rank* of a shape. We can slice  $(e_s[e_n : e_n])$  a shape expression or concatenate  $(e_s @ e_s)$  two shape expressions. For example, suppose a PyTea IR variable  $t$  has shape  $(2, 3, 4)$ . Expression  $t[0]$ , which means the first sub-tensor of  $t$  along the first axis, can be represented inside constraints as  $(2, 3, 4)[1 : \text{rank}(t)]$ , or simply  $(3, 4)$ . In case of expression  $t$ 's shape is unknown( $\alpha_s$ ), the shape of a sub-tensor  $t[0]$  will be represented as  $\alpha_s[1 : \text{rank}(\alpha_s)]$ .

**3.2.1 Constraint generation rules for PyTea IR.** To capture Python semantics and PyTorch shape transformations, PyTea follows the static semantics  $(\sigma \vdash E : e, C)$  of PyTea IR. Judgment  $(\sigma \vdash E : e, C)$  means that the PyTea IR expression  $E$  is statically approximated by a symbolic value expression  $e$  under environment  $\sigma$  in case the constraint set  $C (\subseteq \text{Constraint})$  is satisfied. The environment  $\sigma (\in \text{Var} \xrightarrow{\text{fin}} \text{Value Expr})$  is a finite table that maps variables to symbolic value expressions.

The introduction of constraints happens for branch expressions or PyTorch APIs (See Section 3.2.2). The other expressions will collect constraints from their subexpressions. For example, for an add expression  $(E_1 + E_2)$ , see:

$$\frac{\sigma \vdash E_1 : e_n, C_1 \quad \sigma \vdash E_2 : e'_n, C_2}{\sigma \vdash E_1 + E_2 : e_n + e'_n, C_1 \cup C_2}$$

The result value is symbolically  $(e_n + e'_n)$  where  $e_n$  and  $e'_n$  are symbolic results of  $E_1$  and  $E_2$  respectively. The result constraint set will be a union of the result constraint sets of  $E_1$  and  $E_2$ .

Every symbolic variable originates from external input, e.g., random function or a dataset. Every expression in the constraints is constructed by these variables and constant values.

**3.2.2 Constraint types.** In order to help the constraint resolution engine Z3 come up with a sensible counter-example that violates the derived constraints, we classify the constraints into two exclusive classes: *soft* and *hard* constraints. For Z3 to generate counter-examples, soft constraints can be violated, while hard constraints should not. Thus hard constraints are, for example, those from branch conditions or about the value range of the input. See Figure 5 again. Python built-in `random.randint` function generates an unknown random variable within a given range  $[0, 1]$ . We mark that bound constraint as a hard constraint. On the other hand, `torch.mm` API demands that two input tensors have to be rank-2  $(x, y)$  tensor and the second dimension ( $y$ -coordinate) of the first tensor have to be equal to the first dimension ( $x$ -coordinate) of the second tensor. This condition can be violated under the shape of the inputs, hence we mark it as a soft constraint.

**Hard constraint generation.** Hard constraints are those for inputs and branch conditions. Input conditions restrict the initial ranges of each input. Branch conditions split each path into two.

Consider the following rule.

$$\frac{\begin{array}{l} c_1 = (1 \leq \alpha_n \leq 4) \quad (\text{new } \alpha_n) \\ c_2 = (0 < \alpha'_n) \quad (\text{new } \alpha'_n) \\ c_3 = (0 < \alpha''_n) \quad (\text{new } \alpha''_n) \\ e_s = (\alpha_n, \alpha'_n, \alpha''_n) \end{array}}{\sigma \vdash \text{readImage} : e_s, \{c_1, c_2, c_3\}}$$

The `readImage` API is an image fetching API that creates a new 3-rank tensor which represents color channels, height, and width. The range of color channels is from 1 to 4, i.e., monochrome to RGBA, hence the constraint  $c_1$  in the above rule. The symbolic value is a tensor of shape  $(\alpha_n, \alpha'_n, \alpha''_n)$ .

As another case, consider the following rule.

$$\frac{\begin{array}{l} \sigma \vdash E_1 : e_1, C_1 \quad \sigma \vdash E_2 : e_2, C_2 \\ c = (e_1 \leq \alpha_n \leq e_2) \quad (\text{new } \alpha_n) \end{array}}{\sigma \vdash \text{randInt } E_1 E_2 : \alpha_n, C_1 \cup C_2 \cup \{c\}}$$

The `randInt` API generates a new random variable which is bound to given two numbers. This expression is used from the Python API `random.randint`.

For branching case, see below:

$$\frac{\sigma \vdash E_1 : e_b, C_1 \quad \sigma \vdash E_2 : e, C_2}{\sigma \vdash \text{if } E_1 E_2 E_3 : e, C_1 \cup C_2 \cup \{e_b\}}$$

$$\frac{\sigma \vdash E_1 : e_b, C_1 \quad \sigma \vdash E_3 : e, C_3}{\sigma \vdash \text{if } E_1 E_2 E_3 : e, C_1 \cup C_3 \cup \{\neg e_b\}}$$

The `if` expression creates two paths depending on the branch condition  $e_b$ . If the branch condition can be evaluated to a constant boolean, we can safely drop one branch.

*Soft constraint generation.* Soft constraints are the conditions with which PyTorch APIs must comply for them to run without a shape error. For instance, two operands of a matrix multiplication have to share the same middle dimension, and the reshape operation requires that the number of elements of the input tensor must be matched with the number of elements of the target shape. Each PyTorch API holds unique requirements of input conditions, and PyTea collects these requirements as soft constraints.

Following three rules, for example, PyTea collects such constraints from three representative APIs (`mm`, `reshape` and `transpose`):

$$\frac{\sigma \vdash E_1 : e_s, C_1 \quad \sigma \vdash E_2 : e'_s, C_2 \quad \text{rank}(e_s) = \text{rank}(e'_s) = 2 \quad e''_s = (e_s[0], e'_s[1]) \quad c = (e_s[1] = e'_s[0])}{\sigma \vdash \text{mm } E_1 E_2 : e''_s, C_1 \cup C_2 \cup \{c\}}$$

The `mm` API calculates a matrix multiplication of two 2-rank matrices. The second dimension of the first matrix must be equal to the first dimension of the second matrix following the basic rules of linear algebra.

The `reshape` API redefines the shape of a tensor. Reshaping a tensor does not change or drop the value of a tensor, so the target shape must have the exactly same number of values as the original shape ( $\prod e_s = \prod e'_s$ ):

$$\frac{\sigma \vdash E_1 : e_s, C_1 \quad \sigma \vdash E_2 : e_n, C_2 \quad \sigma \vdash E_3 : e'_n, C_3 \quad e'_s = (e_n, e'_n) \quad c = (0 < e_n) \wedge (0 < e'_n) \wedge (\prod e_s = \prod e'_s)}{\sigma \vdash \text{reshape } E_1 E_2 E_3 : e'_s, C_1 \cup C_2 \cup C_3 \cup \{c\}}$$

The `transpose` API swaps two dimensions of the tensor  $E_1$  along the  $E_2$ -axis and  $E_3$ -axis. Unlike the normal 2-rank matrix transposition, `transpose` slices a tensor with  $(E_2, E_3)$ -plane and transposes each matrix on each cross-section:

$$\frac{\sigma \vdash E_1 : e_s, C_1 \quad \sigma \vdash E_2 : e_n, C_2 \quad \sigma \vdash E_3 : e'_n, C_3 \quad e_1 = e_s[0:e_n] @ (e_s[e'_n]) \quad e_2 = e_s[e_n+1:e'_n] @ (e_s[e_n]) \quad e'_s = e_1 @ e_2 @ e_s[e'_n+1:\text{rank}(e_s)] \quad c = (0 \leq e_n < e'_n < \text{rank}(e_s))}{\sigma \vdash \text{transpose } E_1 E_2 E_3 : e'_s, C_1 \cup C_2 \cup C_3 \cup \{c\}}$$

From this rule, we only consider the shape of the result, not the movement of the value inside the tensor.

---

### Algorithm 1: Offline Constraint Check with SMT Solver

---

**Input:**  $H, S$  - logical conjunctions of hard, and soft constraint sets

**Output:** valid, invalid, dontknow, or unreachable

**Function** `analyze( $H, S$ ):`

```

if checkSat( $H$ ) = unsat then
   $\perp$  return unreachable
else if  $S = \emptyset$  then
   $\perp$  return valid
 $v = \text{checkSat}(\neg(H \rightarrow S))$ 
if  $v = \text{unsat}$  then
   $\perp$  return valid
else if  $v = \text{sat}$  then
   $\perp$  return invalid
else return dontknow

```

---

3.2.3 *Handling path explosion.* Splitting execution paths whenever the analyzer encounters a branch can make the analysis cost grow exponentially. We can ignore some of them using the online constraint check, but we cannot for branches that use run-time input values.

However, we can still avoid path split if both paths behave identically in terms of tensor shape. The conservative conditions are as follows:

- (1) Constraints collected from each path are not dependent on the branch condition, and
- (2) Each path has no global side-effect, and
- (3) Two paths' result symbolic values are the same.

PyTea checks the above conditions locally, within the boundary of the `let` expression containing each branch. When PyTea cannot statically decide on any of the three conditions, it safely assumes the conditions do not hold.

Most branches in PyTorch neural network blocks satisfy the above conditions. Typically, network blocks should result in a tensor with a fixed shape that matches with a requirement of the next block or the training target tensor. Those blocks' feed-forward path will be translated into nested `let` blocks with branches that return the same-shaped tensor.

## 3.3 Constraint check

3.3.1 *Online constraint check.* To reduce the number of constraints and paths, our analyzer eagerly simplifies the symbolic expressions and constraints with primitive arithmetics and comparisons. By our eager, online constraint check, the ranges of each symbol can sometimes be known and be used to judge the subsequent constraints. If a branch condition can be simplified into constant true or false, we can trace only a single branch without splitting the path. If a constraint can be simplified to constant false, we can immediately report that the path is unsafe.

3.3.2 *Offline Constraint check.* PyTea feeds the collected constraints of each path to Z3. Algorithm 1 describes how we classify the Z3's result. The final result of PyTea analyzer can be divided into four cases:

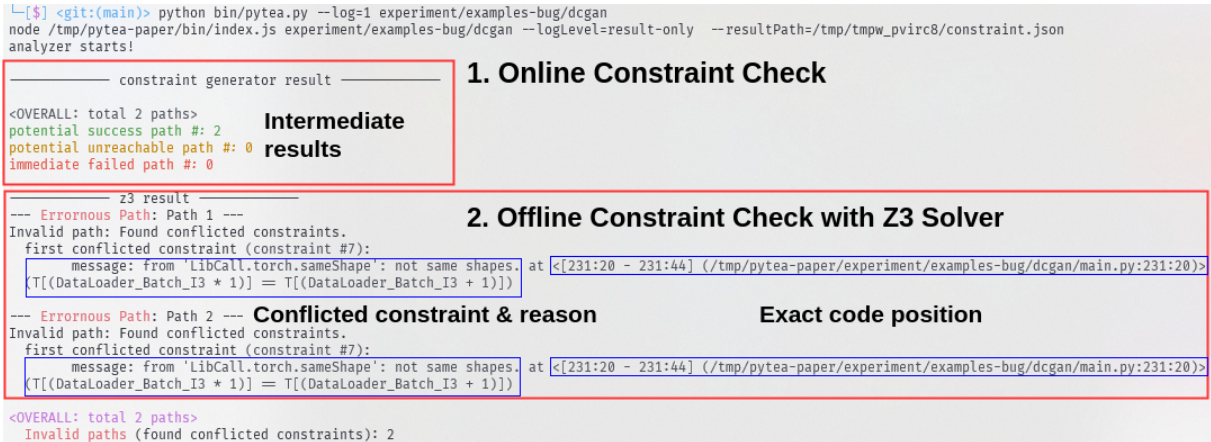


Figure 10: Test result of PyTea command-line tool.

- Valid: Soft constraints are always satisfied under the hard constraints. It guarantees that shape error will not occur from this path.
- Invalid: A possible shape error is detected. There is a counterexample that makes soft constraints false under the hard constraints. We also report the generation position of the first broken constraint.
- Don't know: Z3 failed to decide whether constraints are satisfiable or not.
- Unreachable: There is a conflict between hard constraints in this path. In other words, it is impossible to reach this path under the given conditions. This can happen if a path had passed two contradicted branches.

If every path results in either unreachable or valid path, we can conclude that the input program has no tensor shape error.

## 4 EVALUATION

Our experiments show PyTea's practical performance for real-world applications. To see the practicality of PyTea, we have collected several complete PyTorch applications and shape-related PyTorch bugs. First, we analyzed the official PyTorch example projects from GitHub repository `pytorch/examples`[5]. This repository consists 11 complete PyTorch applications about major machine learning tasks from Generative Adversarial Network (GAN) [24] to Natural Language Processing. We also collected some PyTorch shape mismatch errors from StackOverflow and ran PyTea to statically detect them with PyTea. Finally, we conducted case analyses of several fully-functional, hand-made PyTorch applications such as Stochastic ResNet [21].

*Experiment Settings.* PyTea analyzer is written in mainly TypeScript [2], and communicates with Python scripts to run Z3. We also used Pyright [6] to parse and track Python syntax. The experiments were conducted on R7 5800X CPU, node.js 16.0.0 with TypeScript 4.2.4, and Python 3.8.8 with Z3Py 4.8.10.0. We fixed the epoch size to 1 from the command-line arguments, but used default values for the other settings. We measured the total elapsed time from the

cold boot to the termination of PyTea. The full options and codes are written in the supplementary materials.

*PyTea command-line tool.* Figure 10 shows an example snapshot of the analysis result of the PyTea command-line tool. It has analyzed one of the PyTorch example projects and prints the result of each phase of PyTea. It first prints out the online constraint check results and categorizes each path into three cases, potential success, potential unreachable, and immediate fail. The last one indicates that the online checker has found a constraint that can be false from that path. The potential unreachable path is the path which the online checker has found a false constraint, but there are certain unresolved branch conditions. That path will be checked at the next phase, and PyTea will examine whether the path has conflicted constraints only within the hard constraint set, which means that the path is unreachable from the beginning.

From the second step, PyTea delivers the collected constraint set of each path to Z3 solver and runs the offline constraint checks. The offline check will report the first conflicted constraint and its position of creation, i.e., the exact tensor expression or PyTorch API that causes an error. If the solver does not found any conflicted constraint, PyTea concludes that all the paths are valid, hence no tensor shape error is possible.

### 4.1 Results

*4.1.1 PyTea for PyTorch Examples.* For the experiment, we pass each project twice to the analyzer. For the first pass, PyTea analyzed the main code unmodified, and we check that PyTea does not inform false positives. Then, we injected artificial shape errors, which we subtract one from the first dimension of the target tensor, right before the neural network's loss calculation.

This simple method is decided on purpose. From this experiment, we focused on the speed of PyTea which shows the practicality in order to be integrated to the code editor such as VSCode. This configuration can check the analysis time of the main network, and also confirm that PyTea tracks the tensor operations from the main network thoroughly, and we check PyTea does not report false negative results.

**Table 1: Analysis result of pytorch/examples code repository. The lines of library APIs encapsulated with the analyzer were counted separately. ○: Analysis succeeded and found injected errors, △: Analysis succeeded but requires a modification of the main code (e.g., provide explicit input tensor), ×: Failed to analyze.**

Network	LOC (main + lib)	PyTea	Hattori et al. [20]	Total time (s)
dcgan	3714 (214 + 3500)	○	×	1.75
fast_neural_style	4394 (338 + 4056)	○	×	2.40
imagenet	3820 (320 + 3500)	○	×	2.40
mnist	3607 (116 + 3491)	○	×	1.59
mnist_hogwild	3620 (129 + 3491)	○	△	1.94
reinforcement_learning	180 (180 + -)	×	×	-
super_resolution	3886 (193 + 3693)	△	△	1.57
sqli	223 (223 + -)	×	×	-
time_sequence_prediction	3333 (88 + 3245)	△	×	1.88
vae	3593 (102 + 3491)	○	△	1.70
word_language_model	3278 (361 + 2912)	△	×	1.81

**Table 2: Analysis result of the StackOverflow questions. The numbers in parenthesis denote the URL id of each question.**

Question	PyTea	Hattori et al. [20]
Case 1 (66995380)	○	×
Case 2 (60121107)	○	×
Case 3 (55124407)	○	×
Case 4 (62157890)	○	×
Case 5 (59108988)	○	×
Case 6 (57534072)	○	×

We have compared PyTea against another PyTorch analyzer of Hattori et al. [20]. Table 1 shows the overall results. Among the 11 projects, PyTea successfully analyzed 6 projects without any modification of the original source code. For three projects with a complex data preprocessing stage, PyTea needs a bypass (i.e., code modification) of that stage to infer the shapes of input tensors. PyTea has also succeeded in finding these injected errors. As these results show, PyTea is quick and effective enough to be integrated into code editors. Meanwhile, Hattori et al.’s analyzer failed for almost all benchmarks. Furthermore, since their semi-static approach requires an explicit shape of the input tensor, we needed to feed them an exact network model and input tensors to compare its performance with PyTea.

Although we have aimed to analyze the codes without any modification, two projects are heavily dependent on third-party data managing libraries like OpenAI-Gym [3]. Because, at the moment, we are focusing on the analysis of PyTorch-centered applications, we decided not to support those libraries for now. Supporting more libraries is straightforward and is our future work.

**4.1.2 PyTea for StackOverflow questions.** To show that PyTea can identify yet another set of real-world shape mismatches, we collected some PyTorch shape errors from StackOverflow questions. Recent TensorFlow analyzers [22, 26] used a TensorFlow error dataset collected by Zhang et al. [27], but we manually gathered PyTorch shape mismatch cases rather than using their dataset, because

```

1 class LSTM(nn.Module):
2     def __init__(self, ...):
3         # 7 lines...
4     def forward(self, tokens):
5         # 5 lines ...
6         return out_scores
7
8 model = LSTM(embedding_matrix=np.zeros((1181, 100)))
9 loss_function = nn.NLLLoss()
10 optimizer = optim.Adam(model.parameters())
11
12 ## CUSTOM INPUT
13 input = torch.ones(256, 4, dtype=torch.long)
14 target = torch.ones(256, 4, dtype=torch.long)
15 output = model(input)
16
17 ## ORIGINAL
18 # output: [256 x 4 x 1181], target: [256 x 4]
19 # SHAPE MISMATCH: [256 x 1181] != [256 x 4]
20 loss = loss_function(output, target)
21
22 ## FIXED
23 # output: [1024 x 1181], target: [1024]
24 loss = loss_function(output.reshape(256*4, 1181),
                       target.reshape(256*4))

```

**Figure 11: Example code of StackOverflow question. (Case 2)**

of the fundamental difference of the structures between TensorFlow and PyTorch. We also considered porting the TensorFlow error dataset into PyTorch codes, but we concluded that the ported codes are fairly old and artificial and do not reflect the standard method to build a PyTorch application.

Table 2 gives the analysis results of the 6 questions that we have collected. PyTea could detect every shape mismatch case from those questions. Following the analysis result, we could find the exact error positions and fix the shape mismatch cases. For example, the main code (Figure 11) of Case 2 does not satisfy the shape conditions for the inputs of NLLoss (line 9). The NLLoss module requires that the shape of the first input tensor without the second dimension is equal to the shape of the second input tensor. PyTea found out that NLLoss could generate a shape error from our experiment.



```

1 def load_image(filename, size=None, scale=None):
2     # POTENTIAL ERROR: channel size can be 1.
3     img = Image.open(filename)
4     # img = Image.open(filename).convert('RGB')
5     # ...
6     return img

```

Figure 12: Insufficient preprocessing of image file.

```

1 def forward(self, x):
2     residual = x
3
4     if self.training:
5         # sample random float value
6         sample = self.m.sample().item()
7
8         ### PATH EXPLOSION
9         if sample > 0:
10            out = self.conv1(x)
11            out = self.bn1(out)
12            out = self.relu1(out)
13            out = self.conv2(out)
14            out = self.bn2(out)
15
16            if self.downsample is not None:
17                residual = self.downsample(x)
18            out = out + residual
19        else:
20            if self.downsample is not None:
21                residual = self.downsample(x)
22            out = residual
23        # ...
24
25    out = self.relu2(out)
26    return out

```

Figure 13: Path explosion in Stochastic ResNet block.

We then fixed the code according to the StackOverflow answer, and PyTea checked that every path became valid.

## 4.2 Discovered Errors in PyTorch Applications

We applied PyTea to several realistic PyTorch applications which contain potential shape errors or path explosion. PyTea-found shape errors include the typical type of shape errors that we introduced at Section 1.3. The complete projects and experiment scripts from this section will be in the supplementary material.

**4.2.1 Detecting insufficient data preprocessing.** We found a potential error at the data preprocessing stage from `fast_neural_style` application of `pytorch/examples` repository. As shown in Figure 12, `Image.open` does not guarantee the loaded image has channel 3, i.e., RGB image. Therefore, any training or inference stage with a monochrome image will fail if we miss the channel converting method like line 4. This error was remained from the initial version and was fixed by the latest commit (a3f28a2) of the preprocessing script.

**4.2.2 Handling path explosion.** For a neural network model which contains a runtime path-explosion, PyTea analyzed it without a

```

1 class NTXentLoss(torch.nn.Module):
2     def __init__(self, batch_size, temperature):
3         super(NTXentLoss, self).__init__()
4         self.batch_size = batch_size
5         # ...
6
7     def forward(self, zis, zjs):
8         batch = self.batch_size
9         repr = torch.cat([zjs, zis], dim=0)
10        sim = self.similarity_function(repr)
11
12        ## zis: [B x N], sim: [2B x 2B]
13        ## CONSTRAINT: -sim.shape[0] <= b <= sim.shape[0]
14        l_pos = torch.diag(sim, b)
15
16        # ...
17        diag = torch.eye(2 * b)
18        l1 = torch.diag(torch.ones(b), -b)
19        l2 = torch.diag(torch.ones(b), b)
20        mask = diag + l1 + l2
21        mask = (1 - mask).type(torch.bool)
22        # 'mask' tensor has (4b^2 - 4b) True values.
23
24        negatives = sim[mask].view(2 * b, -1)
25        # shape of 'negatives': (2b, 2b - 2)
26        # ...
27
28        # ...
29        train_loader = DataLoader(
30            train_dataset,
31            batch_size=256,
32            # drop_last=True, # ERROR
33        )
34        losses = train(net, train_loader)

```

Figure 14: Shape inference which requires the exact values of a tensor.

timeout. The `stochastic-resnet` example uses several deep learning techniques, mainly stochastic depth training [21]. See Figure 13. From this application, the building block of the network contains runtime branches (line 9) that can cause a path explosion. PyTea's path handling algorithm can successfully prune those branches and finishes without timeout. (Caveat: the overall data handling is somewhat hard to follow; we did not automatically reduce the repeat count of the main training loop. We explicitly reduced the length of the dataset (CIFAR-10) with a configuration file (`pyteaconfig.json`), and without modifying the code itself.)

**4.2.3 Handling both regular and residual batch sizes in the training loop.** PyTea considers a residual minibatch in the training loop which leads to a shape error, as we discussed in Section 2. We simplified the SimCLR [7, 12] application to a single PyTorch-only script. From line 4 of Figure 14, the main network class `NTXentLoss` takes an exact batch size to initialize itself. So if we omit `drop_last` parameter that removes the last batch at line 32, the last residual minibatch will lead to a crash if the total data size cannot be divided into the batch size. PyTea finds that the inequality between two batch sizes from line 14 of Figure 14 generates a shape error.

### 4.3 Limitation of PyTea

The main focus of PyTea is the detection of shape errors, so it does not perform general value analysis such as tracking the value of the tensor or array index out-of-bound exception.

If a shape of a tensor is dependent on the value of the other tensor, PyTea can miss a shape error. For instance, the `view` method at line 18 of Figure 14 requires that the element count of an input tensor is divisible by  $2b$ . Tensor masking by a boolean tensor (similarity\_matrix[mask]) returns a 1-D tensor whose length is equal to the number of True of the masking tensor. Although lines 10 to 14 guarantee that the masking tensor has  $4b^2 - 4b$  True, we do not know the view API will succeed since we do not track the exact value of a tensor.

## 5 RELATED WORKS

There is only one work [20] of statically detecting shape mismatch of PyTorch applications. Hattori et al. [20] presented a semi-static analysis of PyTorch applications that requires explicit tensor inputs. Because of the path-insensitive and semi-static approach, their tool is premature to fully statically analyze real-world applications. As shown in Table 1, the performance of their tool is impractical.

For TensorFlow applications, the latest static analyzer is Pythia [22], following the same group’s previous work Ariadne [16]. Pythia is dependent on the Doop framework[10, 18] for Java pointer analysis and the Datalog language. Since Pythia’s target is not Python, their coverage of Python and TensorFlow is still insufficient to handle real-world applications. For example, Pythia cannot analyze integer modular operation and tensor indexing and slicing, as shown in Figure 15. ShapeFlow [26] is a tester, a dynamic analyzer with fake TensorFlow libraries that only track shape transformations. Their dynamic approach achieved better performance and coverage than Ariadne and Pythia, but it requires a reduced dummy dataset to run their tool. It cannot detect a possibility of shape mismatch caused by an untested input dataset.

There are several works to solve the shape mismatch problems[14, 16, 20, 22, 26], but they all have fundamental limitations to analyze PyTorch machine learning applications, such as the lack of support for handling external data, branches, and loops. Also, most of them work on TensorFlow applications.

Static analyses for Python programs have also been reported [14, 17]. Notably, Cruz-Camacho’s thesis[14] contains the shape analysis of NumPy[19] array operators. However, their coverage of Python syntax is restricted that custom function and class declaration are not supported. PyExZ3[9] is a value analyzer for Python language that implemented dynamic symbolic executor with Z3 backend. To port it for a shape mismatch problem needs a sizeable overhaul.

## 6 CONCLUSION AND SIGNIFICANCE

We have developed an automatic static analyzer PyTea that detects tensor-shape mismatch errors in PyTorch’s deep neural network code. Our experiments have shown that PyTea’s performance is practical in reality.

*Significance.* The tensor-shape mismatch error is a critical bug in deep neural net training code, yet hard to statically detect by programmers. We presented a solution to this problem: automatic

static analyzer PyTea whose performance is realistic. The analysis design strikes a balance between the cost, accuracy, and coverage with a focus on the typical program structure of PyTorch deep neural net code base. PyTea is classified as a bug-finder, not a verifier: PyTea can have false positives or false negatives in principle, yet we observed no such cases in our experiments. PyTea is ready for public release.

```

1 import tensorflow as tf
2
3 one = 1
4 four = 1
5 if one == 1:
6     four = 4
7
8 target = tf.ones((4, 5))
9
10 with tf.Session() as sess:
11     t0 = tf.ones((3, 4)) # [3 x 4] * [4 x 5]
12     p0 = tf.matmul(t0, target) # Pass: Correct
13
14     t1 = tf.ones((3, 5)) # [3 x 5] * [4 x 5]
15     p1 = tf.matmul(t1, target) # Error: Correct
16
17     t2 = tf.ones((3, 5 % 2)) # [3 x 2] * [4 x 5]
18     p2 = tf.matmul(t2, target) # Pass: False Negative
19
20     t3 = tf.ones((3, 5))[0] # [5] * [4 x 5]
21     p3 = tf.matmul(t3, target) # Pass: False Negative
22
23     t4 = tf.ones((3, 5))[0:1] # [1 x 5] * [4 x 5]
24     p4 = tf.matmul(t4, target) # Pass: False Negative
25
26     t5 = tf.ones((3, four)) # [3 x 4] * [4 x 5]
27     p5 = tf.matmul(t5, target) # Error: False Positive
28     # ...

```

Figure 15: Basic tensor operations that Pythia [22] fail to analyze correctly.

## ACKNOWLEDGMENT

This work was partially supported by Korea Institute for Information & Communications Technology Promotion (No.2021-0-00059), NAVER CLOVA (No. 0536-20200005) and Supreme Prosecutors’ Office of the Republic of Korea (No. SPO2020A1103DIGITALB, No. SPO2021A1103DIGITALB). This work was also supported by BK21 FOUR Intelligence Computing(Dept. of Computer Science and Engineering, SNU) funded by National Research Foundation of Korea(NRF) (419990214639).

## REFERENCES

- [1] 2010. *Pillow*. Retrieved April 26, 2021 from <https://github.com/python-pillow/Pillow>
- [2] 2012. *TypeScript*. Retrieved April 26, 2021 from <https://www.typescriptlang.org/>
- [3] 2016. *OpenAI Gym*. Retrieved April 26, 2021 from <https://github.com/openai/gym>
- [4] 2016. *Torchvision*. Retrieved April 26, 2021 from <https://github.com/pytorch/vision>
- [5] 2017. *pytorch/examples*. Retrieved April 2, 2021 from <https://github.com/pytorch/examples>
- [6] 2018. *Pyright*. Retrieved December 16, 2020 from <https://github.com/microsoft/pyright>
- [7] 2020. *sthalles/SimCLR*. Retrieved May 1, 2020 from <https://github.com/sthalles/SimCLR/tree/e8a690ae4f4359528cfa6f270a9226e3733b7fa>

- [8] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [9] Thomas Ball and Jakub Daniel. 2015. *Deconstructing Dynamic Symbolic Execution* (proceedings of the 2014 marktoberdorf summer school on dependable software systems engineering, the 2014 marktoberdorf summer school on deep ed.). Technical Report MSR-TR-2015-95. <https://www.microsoft.com/en-us/research/publication/deconstructing-dynamic-symbolic-execution/>
- [10] Martin Bravenboer and Yannis Smaragdakis. 2009. Strictly Declarative Specification of Sophisticated Points-to Analyses. *SIGPLAN Not.* 44, 10 (Oct. 2009), 243–262. <https://doi.org/10.1145/1639949.1640108>
- [11] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. arXiv:2005.14165 [cs.CL] <https://arxiv.org/abs/2005.14165> arXiv:2005.14165.
- [12] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. *CoRR abs/2002.05709* (2020). arXiv:2002.05709 <https://arxiv.org/abs/2002.05709>
- [13] François Chollet et al. 2015. Keras. <https://keras.io>
- [14] Elkin Cruz-Camacho. 2019. *Static Analysis of Python Programs using Abstract Interpretation: An Application to Tensor Shape Analysis*. Master’s thesis, Universidad Nacional de Colombia - Sede Bogotá. <http://bdigital.unal.edu.co/72620/>
- [15] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’08/ETAPS’08)*. Springer-Verlag, Budapest, Hungary, 337–340.
- [16] Julian Dolby, Avraham Shimnar, Allison Allain, and Jenna Reinen. 2018. Ariadne: Analysis for Machine Learning Programs. In *Proceedings of the 2nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages (MAPL 2018)*. Association for Computing Machinery, Philadelphia, PA, USA, 1–10. <https://doi.org/10.1145/3211346.3211349>
- [17] Aymeric Fromherz, Abdelraouf Ouadjaout, and Antoine Miné. 2018. Static Value Analysis of Python Programs by Abstract Interpretation. In *10th International Symposium NASA Formal Methods (NFM 2018)*. Springer, Newport News, VA, USA, 185–202. [https://doi.org/10.1007/978-3-319-77935-5\\_14](https://doi.org/10.1007/978-3-319-77935-5_14)
- [18] Neville Grech and Yannis Smaragdakis. 2017. P/Taint: Unified Points-to and Taint Analysis. In 2017 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications. *Proc. ACM Program. Lang.*, Article 102, 28 pages. <https://doi.org/10.1145/3133926>
- [19] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. Array programming with NumPy. *Nature* 585, 7825 (Sept. 2020), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- [20] Momoko Hattori, Shimpei Sawada, Shinichiro Hamaji, Masahiro Sakai, and Shunsuke Shimizu. 2020. Semi-Static Type, Shape, and Symbolic Shape Inference for Dynamic Computation Graphs. In *Proceedings of the 4th ACM SIGPLAN International Workshop on Machine Learning and Programming Languages (MAPL 2020)*. Association for Computing Machinery, London, UK, 1119. <https://doi.org/10.1145/3394450.3397465>
- [21] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. 2016. Deep Networks with Stochastic Depth. In *14th European Conference on Computer Vision*. Springer, Amsterdam, Netherlands, 646–661.
- [22] Sifis Lagouvardos, Julian Dolby, Neville Grech, Anastasios Antoniadis, and Yannis Smaragdakis. 2020. Static Analysis of Shape in TensorFlow Programs. In *34th European Conference on Object-Oriented Programming (ECOOP 2020)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 15:1–15:29. <https://doi.org/10.4230/LIPIcs.ECOOP.2020.15>
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 8024–8035. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [24] Alec Radford, Luke Metz, and Soumith Chintala. 2016. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *4th International Conference on Learning Representations*. San Juan, Puerto Rico. <http://arxiv.org/abs/1511.06434>
- [25] Xavier Rival and Kwangkeun Yi. 2020. *Introduction to Static Analysis: An Abstract Interpretation Perspective*. The MIT Press, Cambridge, MA, USA.
- [26] Sahil Verma and Zhendong Su. 2020. ShapeFlow: Dynamic Shape Interpreter for TensorFlow. arXiv:2011.13452 [cs.LG] <https://arxiv.org/abs/2011.13452> arXiv:2011.13452.
- [27] Yuhao Zhang, Yifan Chen, Shing-Chi Cheung, Yingfei Xiong, and Lu Zhang. 2018. An Empirical Study on TensorFlow Program Bugs. In *Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA 2018)*. Association for Computing Machinery, Amsterdam, Netherlands, 129–140. <https://doi.org/10.1145/3213846.3213866>
- [28] Barret Zoph and Quoc V. Le. 2017. Neural Architecture Search with Reinforcement Learning. In *5th International Conference on Learning Representations*. OpenReview.net, Toulon, France. <https://openreview.net/forum?id=r1Ue8Hxcg>