

# JANUS: Fast and Flexible Deep Learning via Symbolic Graph Execution of Imperative Programs\*

Eunji Jeong  
ejjeong@snu.ac.kr  
Seoul National University

Joo Seong Jeong  
joosjeong@snu.ac.kr  
Seoul National University

Sungwoo Cho  
sungwoocho@snu.ac.kr  
Seoul National University

DongJin Shin  
dongjin.shin@spl.snu.ac.kr  
Seoul National University

Gyeong-In Yu  
gyeongin@snu.ac.kr  
Seoul National University

Byung-Gon Chun  
bgchun@snu.ac.kr  
Seoul National University

## Abstract

The rapid evolution of deep neural networks is demanding deep learning (DL) frameworks not only to satisfy the traditional requirement of quickly executing large computations, but also to support straightforward programming models for quickly implementing and experimenting with complex network structures. However, existing frameworks fail to excel in both departments simultaneously, leading to diverged efforts for optimizing performance and improving usability.

This paper presents JANUS, a system that combines the advantages from both sides by transparently converting an imperative DL program written in Python, the de-facto scripting language for DL, into an efficiently executable symbolic dataflow graph. JANUS can convert various dynamic features of Python, including dynamic control flow, dynamic types, and impure functions, into elements of a symbolic dataflow graph. Experiments demonstrate that JANUS can achieve fast DL training by exploiting the techniques imposed by symbolic graph-based DL frameworks, while maintaining the simple and flexible programmability of imperative DL frameworks at the same time.

## 1 Introduction

In recent years, deep neural networks have been widely used in various application domains such as computer vision, speech, and natural language processing for their powerful capabilities of extracting abstract features from data. Scientists have created deep learning (DL) frameworks – TensorFlow [6], PyTorch [29], Caffe2 [12], MXNet [8], and many more [42, 40, 27, 37, 25, 14, 38, 7] – to improve the performance of deep neural networks in various jobs and promote the use of deep neural networks in both production and research.

Such DL frameworks can be classified into two distinct families depending on their execution models. One family comprises frameworks that base their execution on symbolic graphs constructed from DL programs. The other family consists of frameworks that directly execute DL programs in an imperative manner.

**Symbolic graph execution.** Frameworks such as TensorFlow [6], Caffe2 [12], and MXNet [8] formulate neural networks as symbolic dataflow graphs. Graph vertices denote the states and operations of a neural network, while graph edges indicate the flow of data between vertices. Operations in the graph are executed as their dependencies are solved, similar to how most dataflow systems process dataflow graphs [10, 19]. The graph representation allows the framework to identify which operations can be run in parallel, and apply various compiler optimization techniques such as common subexpression elimination or constant folding to generate optimized versions of graphs. Moreover, it is easy to process dataflow graphs on accelerator devices or deploy graphs across multiple machines by assigning an operation to the appropriate device or machine.

However, the separation of building a symbolic graph and executing it complicates user experience, because users are not actually running any numerical computations when defining neural networks through the framework interface. Rather, they are constructing graphs that will be executed later through separate functions.

**Imperative program execution.** In contrast, frameworks including PyTorch [29], TensorFlow Eager [37], and MXNet Imperative [24] have adopted the execution model of running operations imperatively, without going through a separate graph construction phase. Stemming from popular Python libraries for scientific, numerical computation such as NumPy [41] and Scikit-learn [30], this imperative approach is useful for rapidly experi-

---

\* Preprint version, to appear at NSDI '19

menting and working with new neural network models, particularly those with complex structures. The native control flow statements of Python can be exploited to build models of interest. Unfortunately, skipping the formation of a dataflow graph means that such frameworks lose the chance to apply the many optimizations that were possible in the symbolic graph execution model, leading to significant performance differences for certain neural network models.

The different characteristics of DL frameworks suggest that we cannot achieve high performance and good usability at the same time. To reach high performance, we must sacrifice framework usability to a certain extent, and vice versa. Otherwise, users are forced to resort to an awkward approach of learning how to use several frameworks and switching between them according to the current task in hand.

**From imperative programs to symbolic graphs.** In this paper, we propose to transparently convert imperative Python DL programs into symbolic dataflow graphs before execution. By not altering the user-facing interface for building neural networks, we maintain the flexible programmability of frameworks with imperative execution models. At the same time, behind the scenes, we execute the symbolic graph versions of the imperative programs to enjoy the performance optimizations done by symbolic graph execution models.

However, this approach introduces a technical challenge of capturing the dynamic semantics of an imperative Python program in a symbolic dataflow graph. The dynamic aspects of Python, including dynamic control flow, dynamic typing, and impure functions, must be embedded in a symbolic graph correctly while providing the performance of symbolic graph execution frameworks.

To this end, we present JANUS, a DL framework that achieves the best of both worlds by receiving an imperative DL program as input and creating symbolic graphs of the program accordingly with speculative program context assumptions if possible. JANUS makes environment assumptions on the program context (e.g., constant variables and conditional branches) based on past iterations to simplify the dynamic nature of the program and transform the program into a symbolic graph. These assumptions are speculative, because the context may change during execution; an incorrect assumption results in an invalidation of a symbolic graph, in which case JANUS falls back to imperative execution to guarantee correctness. For design (Section 4.3.1) and implementational (Section 4.3.2) reasons, JANUS converts only the subset of Python programs into the efficient symbolic graphs, but the rest of them still can be executed imper-

---

```

1  class RNNModel(object):
2      def __call__(self, sequence):
3          state = self.state
4          outputs = []
5          for item in sequence:
6              state = rnn_cell(state, item)
7              outputs += [state]
8          self.state = state
9          return compute_loss(outputs)
10
11 for sequence in sequences:
12     optimize(lambda: model(sequence))

```

---

Figure 1: A Python program that implements training process of a recurrent neural network (RNN) in an imperative manner. For each `item` in the sequence, `rnn_cell` function is called to produce the next state required for the next `rnn_cell` invocation. After finishing up processing the whole sequence, the model holds the final `state` by replacing `self.state` attribute for processing the next sequence.

atively, ensuring the full Python coverage.

We have implemented JANUS on TensorFlow 1.8.0 [6]. To demonstrate the performance of JANUS, we evaluated JANUS with 7 imperative DL programs including convolutional, recurrent, and recursive neural networks that extensively use the dynamic features of Python. JANUS converted the programs into symbolic dataflow graphs successfully, trained the models to reach target accuracy with up to 18.7 times higher throughput compared to TensorFlow Eager, while executing the identical imperative programs.

## 2 Challenges and Proposed Solution

### 2.1 Challenges

Converting an imperative program written in Python into a DL dataflow graph brings on many challenges, because dataflow graphs consist of a restrictive set of operations, lacking the dynamic semantics of the programming language. More specifically, various characteristics of a Python program, such as the execution count and execution order of statements, the types of expressions, or the global program execution state, can only be determined after the program is actually executed. For the rest of this paper, we will refer to these characteristics as the *dynamic* features of Python. In contrast, DL dataflow graphs are expected to be defined before the computation starts, to apply aggressive graph optimizations and efficiently schedule the graph operations by viewing the

Frameworks	Imp. pgm	Sym. exec	Correctness			Optimization w/ runtime info	Language
			DCF	DT	IF		
<b>Symbolic:</b> TensorFlow(TF), Caffe2, MXNet	×	○	–	–	–	–	Python
<b>Imperative:</b> PyTorch (PTH), TF Eager, DyNet	○	×	–	–	–	–	Python
<b>Imperative to Symbolic</b>							
Tracing: TF <code>defun</code> , PTH JIT, MXNet Gluon	○	○	×	△	×	○ (unsafe)	Python
JAX	○	○	○	△	×	○ (unsafe)	Python subset
Swift for TensorFlow (S4TF)	○	○	○	–	○	×	Swift
JANUS	○	○	○	○	○	○	Python

Table 1: Comparison of DL frameworks with respect to correctly supported features for converting imperative programs into symbolic graphs ("Correctness") and the ability to optimize the generated graphs with the information given only at program runtime ("Optimization w/ runtime info"). The host language is also specified.

entire graph. In this sense, DL dataflow graphs are usually considered to be *static* [23, 27, 29]. The difference in characteristics makes it difficult to embed dynamic Python features in static dataflow graphs.

Figure 1 depicts a DL program written in Python, of which semantics are difficult to be captured in a dataflow graph **correctly** due to the following representative dynamic features of Python.

- **Dynamic control flow (DCF):** Conditional branches and iterative loop constructs have different execution paths depending on intermediate values. Lines 5-7 of Figure 1 show an example of an iterative loop construct used in a DL program. Such control flow statements are intensively used in Python and must be correctly represented in the dataflow graph.
- **Dynamic types (DT):** Python is a dynamically-typed language, i.e., the type of a Python expression can only be determined at program execution time. The example program in Figure 1 does not have any type annotations (e.g. `int` or `float`), which makes it difficult to statically decide the type of target dataflow graph operations. Furthermore, various non-numerical types of Python, such as lists, dictionaries, and arbitrary class instances, are even harder to be converted into elements of a dataflow graph, of which vertices usually output numerical arrays.
- **Impure<sup>1</sup> functions (IF):** Another useful feature for using Python is the ease of accessing and mutating global states within functions. In Figure 1, the function `__call__` reads from and writes to an object attribute<sup>2</sup> at Lines 3 and 8, to pass the final state of a sequence to the next sequence. Since the modified global states can make the following function call be-

<sup>1</sup>A *pure function* is a function whose return value is determined only by its parameters, and has no side effects.

<sup>2</sup>"class members" in C++ terminology

have differently, such reads and writes of global states must be handled correctly while generating dataflow graphs.

Moreover, correctness is not the only issue when converting an imperative program; achieving the high **performance** of state-of-the-art symbolic graph execution DL frameworks is also a challenge on its own. State-of-the-art frameworks require additional information on dynamic types and control flow in order to optimize graph execution. However, a naïve, one-shot converter would be unable to extract this information from an imperative program before execution, and thus is incapable of supplying frameworks with such hints. For instance, if the input sequence at Line 2 in Figure 1 is expected to always have a fixed length, then that information can be exploited to unroll the following loop at Line 5 when generating the corresponding dataflow graph. It is unclear how a naïve converter would do this without actually executing the program to check the loop length.

## 2.2 Related Works

Concurrent works that try to translate a Python DL program into a dataflow graph either fail to capture the important dynamic semantics of Python, or run in slower performance due to the lack of sufficient information at graph build time. Table 1 summarizes state-of-the-art DL frameworks alongside their execution models and their status regarding the coverage and efficiency of graph conversion support.

Tracing-based graph generation approaches such as PyTorch’s JIT compiler [29], MXNet Gluon [25], and the `defun` [39] functionality of TensorFlow Eager [37] execute the imperative program once, and convert the single execution trace directly into a dataflow graph. Though this approach enables generating optimized symbolic graphs with sufficient information gathered from a specific execution trace, it fails to capture dynamic seman-

tics of the Python interpreter correctly, leading to incorrect computation results for dynamically changing execution path, dynamic types of non-tensor or non-input expressions, or impure functions of Python at runtime. Moreover, these approaches currently do not give any feedback about incorrectly-converted control flows to users, making the problem even worse.

On the other hand, there exist other approaches that select a less-dynamic host language and therefore succeed in capturing the wider semantics of source programs. JAX [14] limits the Python syntax and supports converting only pure-and-statically-composed functions. S4TF [38] supports Swift, losing the merit of supporting Python, the de-facto standard programming language for DL programming, and introduces new programming models that most DL researchers are unfamiliar with. Moreover, since the graph conversion occurs before actually executing the program, these approaches can miss the opportunity to further optimize the graph with the information only obtainable during the program execution. For example, always converting a Python loop into control flow operations can be sub-optimal if the loop iteration count is known to be fixed.

### 2.3 Proposed Solution: Speculative Graph Generation and Execution

Existing optimizers and compilers for dynamic languages suggest a useful technique for performing such conversions from imperative programs to symbolic dataflow graphs: *speculative optimization*. Managed language runtimes have succeeded in exploiting the inherent static nature of dynamic programs which rarely changes during the execution to convert them into static, low-level representations while maintaining correctness. For example, JavaScript just-in-time (JIT) compilers convert dynamic JavaScript programs into efficient machine code, and this conversion is done speculatively assuming that the program inherently maintains some statically fixed structures over repeated executions. In case this assumption breaks, the function falls back to the interpreter and attempts to compile the program again with different assumptions.

We propose to adopt this concept of speculative optimization when converting imperative DL programs into symbolic dataflow graphs. Converting various dynamic features like dynamic control flow and impure functions correctly may impose some inevitable overheads if we generate dataflow graphs in a conservative manner. To overcome this challenge, JANUS makes assumptions about the program’s behavior based on the runtime profiling information, and generates a symbolic graph tai-

lored for the assumptions. This speculatively constructed dataflow graph can show much better performance compared to the conservative counterpart due to specializations. If the assumptions do not hold, JANUS builds a new dataflow graph based on different assumptions, just like JavaScript JIT compilers. Since a DL program comprises a number of iterations of an optimization procedure, the speculative approach is a good fit since the interpreter is likely to execute specific code blocks of the program repeatedly.

Unlike the JIT compilers of managed language runtimes, however, the goal of JANUS is not to optimize the host language execution itself. In fact, when running imperative DL programs, the execution time of the language runtime is usually much shorter compared to the execution time of the mathematical operations for DL, such as convolution or matrix multiplication. However, since these mathematical operations are usually implemented in separate low-level language like C++, existing JIT compilers of managed language runtimes would execute them just as separated function invocations. Under such execution model, it is impossible to see the multiple mathematical operations at once and apply compiler optimizations or execute them in parallel. On the other hand, JANUS understands the function invocations for such mathematical operations, and converts them into appropriate target graph operations, which can be optimized and be executed efficiently by symbolic graph executors.

## 3 JANUS System Design

In this section, we introduce JANUS, a DL framework that receives an imperative DL program and either executes it as is directly, or generates a symbolic graph version of the program and executes the graph instead. JANUS assumes that the input program is written using the API and the programming model of existing imperative DL frameworks like TensorFlow Eager [37]. From the user’s point of view, the whole graph conversion and execution process is done transparently; in other words, the given DL program is automatically transformed into a corresponding graph representation without any interactions with the user.

Figure 2 depicts the system components and the overall execution model of JANUS. The common case in which an efficient dataflow graph is utilized is depicted as a solid line in the figure, while the rare case where the graph representation is not available is depicted as a dotted line.

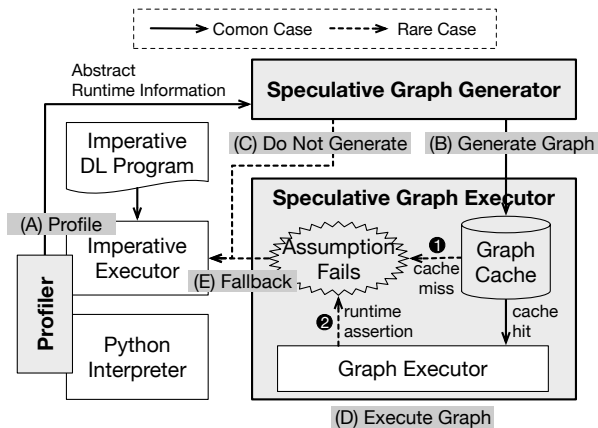


Figure 2: An illustration of the execution model of JANUS, showing how a DL program is processed by several components. *Profiler* observes imperative program execution and collects information to make the realistic assumptions. *Speculative Graph Generator* generates dataflow graphs from the program and hands the optimized graphs over to *Speculative Graph Executor*. The *Speculative Graph Executor* actually runs the generated graph and handles assumption failures.

### 3.1 Fast Path for Common Cases

**Runtime profiling.** Once JANUS receives a DL program, the program is first executed imperatively, while the *Profiler* gathers abstract runtime information required for making reasonable assumptions (Figure 2 (A)). Various information is collected, including control flow decisions on conditional branches, loop iteration counts for iterative loop constructs, variable type information, non-local variables, object attributes, and so on.

**Symbolic graph generation.** After a sufficient amount of information has been collected, the *Speculative Graph Generator* tries to convert the program into a symbolic dataflow graph with the assumptions based on the runtime information (Figure 2 (B)). To avoid making any hasty generalizations, JANUS does not begin graph generation until the executor has profiled the program for a certain amount of iterations.<sup>3</sup> First, JANUS traverses the abstract syntax tree (AST) of the DL program and generates the corresponding graph elements for each AST node, along with assertion operations that can validate the context assumption at runtime. Since JANUS targets DL programs, operations for automatic differentiation and model parameter updates are also automatically inserted if necessary. Next, the generated

<sup>3</sup>We found that 3 iterations were enough to come up with a decent program context assumption, for our experimental workloads.

graph is further optimized by the post-processor, of which optimizations were not applicable to the original imperative DL program. Finally, the optimized graph and the assumption that were used to generate the graph are saved into the graph cache.

**Graph execution.** If a graph representation with correct assumptions regarding the program context is available, the *Speculative Graph Executor* executes the symbolic graph (Figure 2 (D)). Note that the same graph can be reused multiple times, given that the runtime context assumption holds for future invocations.

### 3.2 Accurate Path for Rare Cases

**Assumption failure.** Handling the assumptions is important to guarantee the correctness of the converted graph. If an assumption is proven to be wrong, the associated graph cannot be executed for the current runtime as it may produce incorrect results. Instead, JANUS falls back to the imperative executor (Figure 2 (E)) and resumes runtime profiling to make more relaxed assumptions for subsequent executions.

Assumptions that can be validated before actually executing the associated graph, such as type assumptions on input arguments, are checked when retrieving the graph from the graph cache (Figure 2 ①). In the unfortunate case where such an assumption is wrong, JANUS regards this as a cache miss and falls back to imperative execution.

On the other hand, for assumptions that can only be validated during graph execution (Figure 2 ②), JANUS cannot simply abort the current execution to fall back to the imperative executor, because the global state may have been changed during the current execution. To solve this issue, JANUS defers state update operations until the very end of the computation (Section 4.2.3). This way, even if an assumption turns out to be wrong during computation, no state update operation has been triggered yet and thus no state has been mutated. Knowing this, the system can safely stop the current execution. States are updated only when every assumption is checked and the computation finishes, in an "all-or-nothing" manner.

In order to validate an assumption, a runtime assertion is encoded into the symbolic graph as an operation called `AssertOp`. The `AssertOp` aborts the graph execution if the given condition fails. It also reports which assumption has been broken, and this information is used to give up further optimizations that rely on the assumptions that repeatedly break.

**Imperatively executed programs.** With Turing-complete graph representations, any Python program

can be represented as a symbolic graph, in theory. However, the *Speculative Graph Generator* does not convert every single Python feature into a symbolic graph operation. For example, to ensure the "all-or-nothing" characteristic of state updates, programs that include invisible state mutations are not converted into symbolic graphs. Some complicated Python features such as coroutines and generators are also not converted, since they do not have any clear graph representations. Section 4.3 describes the design choices and current limitations of the *Speculative Graph Generator* in terms of Python coverage. In spite of such limitations of the *Speculative Graph Generator*, however, it is worth noting that JANUS users can still freely use the all features of Python on the imperative executor.

## 4 Symbolic Graph Generation

In this section, we describe in detail how JANUS converts an imperative DL program into a symbolic dataflow graph. We start the section by showing the conversion process of a basic DL program free of dynamic features (Section 4.1). Next, we explain how JANUS converts dynamic features of Python, including dynamic control flow, dynamic types, and impure functions, into symbolic graph operations. (Section 4.2). JANUS uses the abstract runtime information collected at runtime to simplify the dynamic program and treat it as a program of only static aspects, which is then easily transformed into a static graph. Finally, we discuss the Python coverage limitations of the *Symbolic Graph Generator* (Section 4.3). More thorough discussion about the Python coverage of JANUS is in Appendix A.

For simplicity, we describe our design using various operations of TensorFlow [6], a widely-used DL framework. However, our design is not necessarily coupled with TensorFlow and can be applied to other DL frameworks such as MXNet.

### 4.1 Graph Generation Basics

Figure 3a is a simple, imperative Python program that calculates a linear model, written as a pure function without any dynamic control flow or arbitrary Python objects. We use this program as an example to show the basic graph conversion process.

Input parameters ( $x$  and  $y$ ) are converted into graph input objects that require external inputs in order to execute the graph. In the case of TensorFlow, this corresponds to the `Placeholder`<sup>4</sup> object. At runtime, they are filled

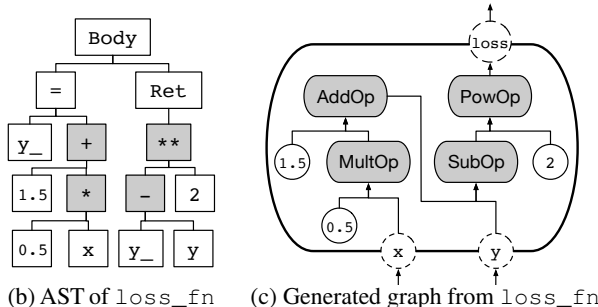
<sup>4</sup>Placeholders are unique operations that generate errors unless they are provided with external inputs before graph execution. TensorFlow expects users to feed a `Placeholder` dictionary `{ph1: v1, ph2: v2, ...}` to the system.

```

1 def loss_fn(x, y):
2     y_ = 0.5 * x + 1.5
3     return (y_ - y) ** 2

```

(a) Source code of a DL program calculating a linear model



(b) AST of `loss_fn` (c) Generated graph from `loss_fn`

Figure 3: The Python source code, AST, and symbolic graph of a simple linear model that receives several external inputs. The static features of the program are represented as nodes in the AST, which in turn are converted to vertices of the symbolic graph.

with the actual argument values. The return value of the return statement is marked as the computation target of the graph, so that we can retrieve the value after executing the graph.

Python literals such as `0.5`, `1.5` and `2` are simply converted into operations that output constant values – `ConstantOp` for TensorFlow. The conversion of mathematical operators is done by finding the corresponding mathematical graph operations and replacing them one-to-one. For standard Python operators such as `+` and `**`, JANUS places the appropriate primitive calculation operations in the graph, like `AddOp` and `PowOp` for TensorFlow.

An assignment to a Python local variable and a value retrieval from the same variable is converted into a connection between two operations, just as in Pydron [26]. Figures 3b and 3c show how such a connection is made for the program in Figure 3a, along with the rest of the program.

### 4.2 Dynamic Features

In addition to the basic features, JANUS converts the dynamic features of Python into the elements of the symbolic DL graph as well to provide the performance of dataflow graphs while maintaining the same programmability of imperative DL frameworks. Moreover, JANUS exploits the fact that the dynamism in Python DL programs can often be simplified to static dataflow, treating a dynamic program as a program of only static aspects with appropriate program context assumptions.

Context assumptions are generated based on the profile information JANUS profiles at runtime.

#### 4.2.1 Dynamic Control Flow

**Basic translation rules.** Among various dynamic control flow statements, JANUS focuses on conditional branches, loop constructs, and function calls, similar to Pydron [26]. As shown in Pydron, these three constructs are enough to express most complex dynamic control flows in Python. Furthermore, they can all be expressed using special control flow graph operations proposed in recent works [43, 20] as follows.

Python’s conditional statement, the `if` statement, can be obtained by combining *switch* and *merge* primitives. The *switch* and *merge* primitives, originating from classic dataflow architectures [11, 9, 28], act as demultiplexers and multiplexers, respectively, selecting a single path to pass their inputs or outputs. In TensorFlow, the `SwitchOp` and `MergeOp` [43] operations serve as symbolic dataflow graph counterparts for these primitives, allowing JANUS to plant conditional branches in graphs.

The iterative statements of Python, `while` and `for`, are handled by using the *switch* and *merge* primitives together with loop context primitives that hold iteration *frames*. TensorFlow conveniently provides `NextIterationOps` [43] for tracking iterations and passing iteration frames.

Finally, for function calls, a separate graph is generated for the callee function, and a function invocation operation that points to the generated graph is inserted in the position of the function calls. Recent work proposes a TensorFlow implementation of this operation called `InvokeOp` [20], which can even represent an invocation of a recursive function.

#### **Speculative graph generation: unrolling and inlining.**

If JANUS detects that only a single particular path is taken for a certain control flow operation during profiling, JANUS presumes that the control flow decision is actually fixed. The system replaces the control flow operation with a *assertion* operation that double-checks the assumption for this control flow decision, and proceeds with graph generation as if the control flow operation were unrolled. This allows JANUS to remove control flow operation overheads and apply graph optimizations such as common subexpression elimination or constant folding in broader portions of the graph. If the *assertion* operation fails, JANUS falls back to imperative execution.

To be more specific, for conditional branches, if the program takes only one side of the branch during profiling, JANUS generates that particular side of the branch

in the final graph without any *switch* or *merge* primitives and adds an *assertion* operation that can detect a jump to the other side of the branch. For iterative statements, if the number of iterations of a loop is discovered to be fixed, JANUS unrolls the loop with this fixed iteration count, and adds an *assertion* operation to check that the number of iterations is indeed correct.

For function calls, if the callee is expected to be fixed for a function call at a certain position, JANUS inlines the callee function body inside the caller unless that function call is identified as a recursive one. In addition, for callee functions whose implementation is already known for JANUS, e.g., the functions provided by the framework such as `matmul()` or `conv2d()`, or Python built-in functions like `print()` or `len()`, JANUS adds the corresponding graph operations which behave the same as the original callee functions, based on the prior knowledge about their behaviors. Section 4.3.1 includes more details and limitations about such function calls.

#### 4.2.2 Dynamic Type

**Basic translation rules.** The types of all expressions within a Python program must be known before JANUS can convert the program into a symbolic graph, because graph operations require operands to have fixed types. This is a challenging task for Python programs because we cannot determine the type of an arbitrary Python expression before actually executing the expression. Fortunately, it is possible to infer the types of some expressions, given the types of other expressions; for example, it is clear that the variable `c` in `c = a + b` is an integer if `a` and `b` are integers.

As a basic rule, JANUS converts numerical Python values such as scalars, list of numbers, and `numpy` [41] arrays into corresponding tensors, and converts non-numerical values, including arbitrary class instances, into integer-typed scalar tensors which hold pointers to the corresponding Python values. Next, JANUS infers the types of other expressions that are derived from expressions covered by the basic rule.

**Speculative graph generation: specialization.** Expressions whose types cannot be inferred from other expressions require a different measure. For instance, it is impossible to identify the types of input parameters for functions, or Python object attribute accesses (`obj.attr`) without any external clues. Inferring the return types of recursive function calls is also impossible due to the circular dependencies. To make proper assumptions about the types of such expressions, *Profiler* observes the types of the expressions during imperative executions. Given these context assumptions, JANUS can

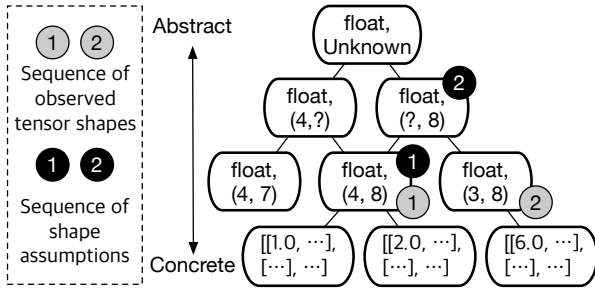


Figure 4: Type, shape, and value specialization hierarchy for an example tensor.

finish inferring the types of remaining expressions, and construct a specialized dataflow graph accordingly.

JANUS makes further assumptions about the expressions to apply more aggressive optimizations. If the profiled Python expressions always have the same values, JANUS converts them into constant nodes in the dataflow graph. Moreover, for numerical expressions, we can try to specialize the shape of tensors before constructing the graph. With statically determined values or shapes, the graph can be optimized even further by techniques such as constant folding or loop unrolling.

Figure 4 shows an example hierarchy of shapes and values that a certain tensor may have. After profiling the first few runs, JANUS finds out that even though the values of the tensor are different every time, they all have the same shape, for example (4, 8), as in the Figure. JANUS exploits this information to generate a dataflow graph with an assumption that the shape of this tensor is (4, 8).

When the assumption fails, JANUS tries to relax the assumption and generate more generalized version. For instance, in case the tensor has a shape (3, 8) for the next iteration to process a different size of mini-batch, JANUS modifies the assumption to suit both shapes (4, 8) and (3, 8), resulting in another dataflow graph with a shape assumption of (?, 8). The system does not have to repeat the graph generation process for a possible future case in which the example tensor has yet another unpredicted shape of (2, 8) or (6, 8).

### 4.2.3 Impure Functions

**Naïve translation rules.** It is common for a Python function to access global variables to calculate return values and have side-effects, mutating its enclosing Python context during execution. Likewise, it is common for a Python DL program to read from and write to global states, such as global and nonlocal variables, or heap objects. JANUS respects this characteristic and handles global state accesses and updates alongside sym-




Figure 5: Symbolic dataflow graph generated graph from Figure 1 and the global states.

bolic graph execution.

A trivial solution is to use TensorFlow’s `PyFuncOps`, which can execute arbitrary Python functions as graph operations. A function for reading and updating a certain global state can be created and inserted in the appropriate position within the graph. However, this trivial approach has clear limitations. First, since only one Python function can be executed at a time due to the global interpreter lock (GIL), the overall performance can be reduced when multiple operations should be executed in parallel. It also complicates the fallback mechanism of JANUS. If a global state has already been mutated before the fallback occurs, instead of starting the imperative executor from the function entrance at fallback, execution must start from the middle of the function to be correct, by mapping the state update operation with the corresponding Python bytecode.

### Optimized graph generation: deferred state update.

To make things simpler and also faster, JANUS does not mutate globally accessible states in place on the fly. JANUS instead creates local copies of global states, and mutates only the local copies during symbolic graph execution.

Figure 5 shows the symbolic dataflow graph version of the program in Figure 1, which includes the object attribute expressions (`self.state`) that access and mutate the global states. We add new graph elements `PyGetAttrOp` and `PySetAttrOp` to represent Python attribute read and write. Each of them receives an object pointer (`0xb84c`) and a name of the attribute (`"state"`) as inputs, and behaves as follows: ① The `PyGetAttrOp` can access Python heap to read the state unless a corresponding local copy exists. ② When the `PySetAttrOp` wants to update the attribute, a new value is inserted to the local copy instead of directly updating the Python heap. ③ Further read and write operations are redirected to the local-copied attribute. Note that JANUS inserts appropriate dependencies between `PyGetAttrOps` and `PySetAttrOps` if necessary to prevent any data hazards. ④ After the graph executor finishes this run, the local copies are written back to the



Python heap. Subscript expressions (`obj[subscr]`) are similarly implemented with equivalent custom operations.

By not mutating the Python heap directly, JANUS can always bypass Python GIL to execute more read and write operations in parallel. Also, the fallback mechanism of JANUS can be simplified thanks to the all-or-nothing based state update mechanism.

### 4.3 Imperative-Only Features

Albeit being able to support a wide range of imperative DL programs, the current JANUS *Symbolic Graph Generator* does not convert some particular features of Python into dataflow graph elements. Programs with such features are only executed on the imperative executor.

#### 4.3.1 Coverage Limitations from Design

**Alignment with the design principles.** To be aligned with the design of JANUS in previous sections, the JANUS graph generator does not convert some features of Python. For example, to keep the implementation of local copies of global state simple (Section 4.2.3), Python object with custom accessor functions (i.e., `__getattr__` or `__setattr__`, ...) are not supported by the JANUS graph executor. Also, a function should always return the same type of value, to infer the type of call expression (Section 4.2.2).

**External function calls.** As described in Section 4.2.1, JANUS must understand the behavior of functions to convert them into corresponding graph operations. In this sense, the JANUS graph generator converts the external functions, i.e., the framework-provided functions or foreign functions<sup>5</sup>, into dataflow graph operations based on a separate white-list. Most of the framework-provided functions such as `matmul` or `conv2d`, and many commonly-used Python built-in functions such as `print` or `len` are included in this white-list. We plan to cover more functions in the standard Python library.

JANUS handles such external functions with extra caution to ensure correctness. First, since the underlying assumption here is that the implementation of external functions never changes, JANUS prohibits the modification of the functions included in the whitelist. Also, if a framework (or foreign) function includes state mutation (e.g., `assign()` in TensorFlow), the execution of the corresponding graph operation is deferred until all the other assumptions are validated, under the same principle about the deferred state update in Section 4.2.3.

---

<sup>5</sup>functions written in the languages other than Python

#### 4.3.2 Coverage Limitations from Implementation

Currently, JANUS does not cover a few features from Python that do not have clear graph representations. Such Python features include *coroutines*, *generators*, and in-line class definitions and in-line import statements. We plan to support these features as future work.

## 5 Implementation

JANUS is implemented on top of TensorFlow [6] 1.8.0 and CPython [13] 3.5.2. JANUS exploits the existing TensorFlow graph executor and TensorFlow Eager imperative executor as the components. In this section, we explain the modifications to existing systems, and then describe how we implemented data-parallel training.

**Modifications to existing systems** TensorFlow has been modified for several reasons. First, to transparently separate out the neural network computation from the entire Python program without extra user intervention, the automatic differentiation functionality of TensorFlow Eager is modified to trigger JANUS graph conversion. Since the users specify the scope of program to be differentiated by the framework, JANUS converts the differentiated target function into the symbolic graph. Second, to share the model parameters between eager mode and graph mode, JANUS slightly modifies the parameter storing mechanism of TensorFlow Eager. Third, several custom operations had been added, including the `InvokeOp` or `PyAttrOp` as described in earlier sections.

CPython has also been modified to have bytecode-level instrumentation functionality for transparent profiling. Without modifying the interpreter, instrumentation for the profiling should exist at the source-code level, which would significantly affect the performance and the debuggability of the imperative execution.

**Data-Parallelization on JANUS** Using multiple machines equipped with multiple GPUs is a common approach for accelerating deep learning jobs. We integrate JANUS with Horovod [32], a distributed training module for TensorFlow that encapsulates the MPI-style [16] communication as an operation inside the symbolic graph. After converting an imperative program into a dataflow graph, JANUS inserts appropriate communication operations to the graph in order to get the average of gradients generated by multiple workers. Since the generated dataflow graph contains both communication and computation operations, we can parallelize communication and computation operations and achieve better performance compared to distributed training with the imperative execution model.

## 6 Evaluation

This section presents experimental results that show how imperatively defined deep learning programs can be executed both correctly and efficiently when converted into symbolic graphs on JANUS.

### 6.1 Experimental Setup.

**Frameworks.** As baseline frameworks representing symbolic graph execution frameworks and imperative execution frameworks respectively, we use TensorFlow [6] and TensorFlow Eager [37]. We could run the same DL program on JANUS as on TensorFlow Eager, thanks to the transparent graph conversion feature of JANUS. In addition, to demonstrate the correctness of graph conversion of JANUS, we also compare JANUS with TensorFlow `defun` [39], which implements a trace-based graph conversion mechanism. TensorFlow-based frameworks have been chosen to avoid implementation-dependent performance differences.

**Applications.** We have evaluated JANUS with three convolutional neural networks (LeNet [22], ResNet50 [17], Inception-V3 [35]), two recurrent neural networks (LSTM [44], LM [21]), and two recursive neural networks (TreeRNN [33], TreeLSTM [36]).

Category	DCF	DT	IF
Conv. NNs	cond.expr.	unk.type	-
Recurrent NNs	loop	unk.type	global var.
Recursive NNs	recursion	unk.type	heap object

Table 2: DL applications and their dynamic features.

These models are implemented in an imperative programming style, using a number of dynamic features in Python as shown in Table 2. First, convolutional neural network models ResNet50 and Inception-V3 have conditional statements for handling batch normalization [18], which make them behave differently under particular conditions when training and evaluating the model. Next, recurrent neural network models include Python `for` loops, and they also include global state mutation statements to retain hidden states inside the models. Finally, recursive neural network models require all three kinds of dynamic features. First, they include recursive function calls, and conditional statements to separate recursion base cases and inductive cases. Second, they include values with undecided type, the return type of a recursive function is unknown until the function returns certain values. Lastly, They include the Python object access to fetch the information of the current subtree. The implementation of TreeLSTM on TensorFlow

follows the recent recursion-based implementation with `InvokeOp` [20], and JANUS can convert the imperative Python program into the similar recursion-based graph. All models use Python function calls, including Python class methods of high-level deep learning programming APIs such as Keras [2]. Training data instances fed into each neural network have different shapes over different training iterations, when the length of the dataset cannot be divided by the batch size.

As the datasets and the mini-batch sizes used when training models, ImageNet [31] was used for ResNet50 and Inception-V3 models, and the Stanford Sentiment Treebank(SST) [34] dataset was used for TreeRNN and TreeLSTM. MNIST [3], Penn Treebank(PTB) [44], and One Billion Word Benchmark(1B) [1] datasets were used for LeNet, LSTM, and LM models, respectively. Batch sizes 50, 64, 64, 20, 256, 25, and 25 per GPU were used for LeNet, ResNet50, Inception-V3, LSTM, LM, TreeRNN, and TreeLSTM models respectively.

**Environments.** All applications are evaluated on a homogeneous GPU cluster of 6 machines, connected via Mellanox ConnectX-4 cards with 100Gbps InfiniBand. Each machine is equipped with two 18-core Intel Xeon E5-2695 @ 2.10 GHz, and 6 NVIDIA GeForce TITAN Xp GPU cards. Ubuntu 16.04, horovod 0.12.1, CUDA 9.0, cuDNN 7, OpenMPI v3.0.0, and NCCL v2.1 are installed for each machine.

LeNet and LSTM-PTB models are evaluated on a single GPU, since these models and the datasets are regarded to be too small to amortize the communication cost of parallel execution. Similarly, TreeRNN and TreeLSTM models are evaluated on a single CPU, since these models and datasets are regarded to be too small to amortize the communication between CPU and GPU. The other models are evaluated using multiple GPUs. ResNet50 and Inception-V3 models are evaluated using up to 36 GPUs, and LM is evaluated on up to 12 GPUs. The network bandwidth made the throughput of LM saturated on more than 2 machines with MPI-style communication, due to the huge parameter size of LM (0.83 billion parameters). Therefore, model convergence of LM is experimented with 6GPUs.

### 6.2 Model Convergence

Figure 6 shows how the neural networks converge on various underlying frameworks, with ResNet50 with the ImageNet dataset, LSTM with the 1B dataset and TreeLSTM with the SST dataset on four frameworks. We followed the hyperparameter settings reported in the literatures [15, 4, 5]. For all evaluated models, JANUS, TensorFlow, and TensorFlow Eager succeeded to make the neu-

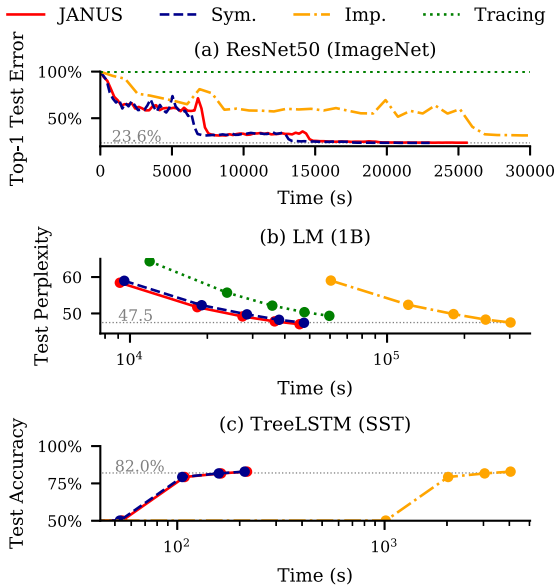


Figure 6: (a) The test error of ResNet50(ImageNet), (b) validation perplexity of LM(1B), and (c) test accuracy of TreeLSTM(SST) models implemented in JANUS, TensorFlow (Sym.), TensorFlow Eager (Imp.), and TensorFlow defun (Tracing). Each marker in LSTM and TreeLSTM convergence graph represents each training epoch, describing that TensorFlow defun does not converge properly after the same number of epochs compared to other frameworks.

ral networks converge correctly as reported in original papers: 23.7% top-1 error for ResNet50-ImageNet, perplexity 47.5 for LM-1B, and 82.0% binary accuracy for TreeLSTM-SST. Also, JANUS could make the model to converge 18.7 times faster than TensorFlow Eager, while executing the identical imperative program. The performance difference between JANUS and TensorFlow was within 4.0%.

However, trace-based TensorFlow defun failed to make the models to converge correctly. The ResNet50 model includes the conditional statement to distinguish the behavior of batch-normalization [18] layer on model training and evaluation. If a user evaluates the initial accuracy before training the model by manipulating the model object attribute, TensorFlow defun converts the first execution trace into graph operations, which silently leads to an inaccurate result. Similarly, the LM model does not converge properly with TensorFlow defun, since it failed to capture state passing across sequences, due to its trace-based conversion mechanism. The TreeLSTM model could not be converted into the symbolic graph at all with TensorFlow defun, since it does not support recursive function call. TensorFlow

Model	(A) Imp.	(B) JANUS	(C) Sym.	$\frac{(B)}{(A)}$	$\frac{(B)}{(C)} - 1$
LeNet	7.94k	25.84k	26.82k	3.25x	-3.6%
ResNet50	188.46	200.37	207.39	1.06x	-3.4%
Inception-V3	108.36	119.32	124.33	1.10x	-4.0%
LSTM	2.75k	22.06k	22.58k	8.03x	-2.3%
LM	19.02k	40.18k	40.45k	2.11x	-0.7%
TreeRNN	20.76	988.72	928.66	47.6x	+6.5%
TreeLSTM	7.51	138.12	141.71	18.39x	-2.5%

Table 3: Training throughput of all models evaluated on a single machine with a single GPU in JANUS, TensorFlow (Sym.), and TensorFlow Eager (Imp.). The numbers represent processed images/s for CNN models, processed words/s for RNN models, and processed sentences/s for Recursive NN models.

Eager converges slowly, since its training throughput is much lower than TensorFlow and JANUS. We next analyze the training throughput of the frameworks, excluding TensorFlow defun, which fails to make models converge correctly.

## 6.3 Training Throughput

### 6.3.1 Single-machine Throughput

Table 3 presents the training throughput of all models executed with JANUS, TensorFlow Eager, and TensorFlow on a single machine with a single GPU. As shown in the table, JANUS outperforms TensorFlow Eager (imperative execution) by up to 47.6 times, and shows throughput similar to TensorFlow (symbolic graph execution) by up to 4.0% performance degradation. JANUS even performs slightly better (+6.5%) for recursive neural networks, since there is no need to pre-process the input sentences, which are the tree-structured Python objects.

JANUS achieves bigger performance gain on recurrent and recursive neural networks than convolutional neural networks since recurrent and recursive neural networks have many concurrently executable operations. In addition, the performance gain of JANUS on a single machine is larger on models with fine-grained graph operations, since the effect of the Python interpreter overhead is more severe when the computation time of an operation is short.

For large convolutional neural networks such as ResNet50 or Inception-V3, optimized GPU kernel computation accounts for most of the computation time, which makes the performance difference among JANUS, TensorFlow, and TensorFlow Eager relatively small.




Figure 8: Training throughput for the ResNet50, Inception-V3, and LM models on JANUS, TensorFlow, TensorFlow Eager, using varying numbers of GPUs. The scale factor of JANUS is 0.77, 0.81, 0.18 for each applications, which is much bigger than the scale factor of TensorFlow Eager (0.24, 0.24, 0.14 each), and similar with the scale factor of TensorFlow (0.81, 0.80, 0.18).




Figure 7: The contribution of optimizations to improve training throughput. Optimizations are cumulative. **IMP**: Imperative, **BASE**: JANUS without following optimizations, **+UNRL**: control flow unrolling, **+SPCN**: type specialization, **+PARL**: graph executor with 72 threads in threadpool (JANUS)

**Optimization Effect.** Figure 7 analyzes the cause of the performance improvement of JANUS in detail. Converting the imperative program into the symbolic graph without any following optimizations (**BASE**) enabled up to 4.9x performance improvement compared to the imperative execution (**IMP**). It removes the Python interpreter and framework code overhead, which has the bigger effect when each graph operation is relatively smaller. Control flow unrolling(**+UNRL**) and type specialization(**+SPCN**) enable more aggressive compiler optimizations. On recurrent neural networks, **+UNRL** improved the performance of LSTM-PTB and LM-1B by 2.09x, 1.04x. The control flow statements in CNNs and Recursive NNs could not be unrolled due to their dynamicity. **+SPCN** enabled some compiler optimizations and improved the throughput up to 18.3% in small neural networks. Finally, executing multiple operations in parallel(**+PARL**) improved the throughput up to 9.81x.

Especially higher gain could be achieved for recursive neural networks, since there exist many operations that could be executed in parallel in multiple independent tree nodes. We have also measured the effect of assumption validation, but the effect was negligible (in the error range), since the `AssertOps` can be executed with the other operations in parallel.

### 6.3.2 Scalability

We evaluated JANUS on the cluster to see if JANUS can exploit data parallelism with multiple machines. Figure 8 shows the scalability of ResNet50, Inception-V3 and LM models on JANUS, TensorFlow, and TensorFlow Eager on the cluster with 36 GPUs(12 GPUs for LM). We measured the scale factor, which is defined as  $Multi-GPU\ Throughput / (Single-GPU\ Throughput \times Number\ of\ GPUs)$ . JANUS achieves similar scalability (scale factor 0.77, 0.81, 0.18 each) as TensorFlow (0.81, 0.80, 0.18 each), but TensorFlow Eager does not scale well (0.24, 0.24, 0.14 each), due its inability to overlap computation and communication.

The slight difference in the scalability of ResNet50 comes from the under-optimized input pipeline of TensorFlow Eager, which JANUS also uses. Optimizing the input processing pipeline for JANUS will further reduce the performance difference between JANUS and TensorFlow. We leave this optimization as a future work.

## 7 Conclusion

In this paper, we introduced JANUS that enables imperatively defining a deep neural network and executing it as a symbolic dataflow graph. JANUS converts the imperative deep learning program written in Python into a dataflow graph by exploiting the inherently static nature of deep learning programs to speculatively generate and execute symbolic dataflow graphs. JANUS solves the

challenges of converting important dynamic features of Python into efficient graph elements. Experimental results showed that JANUS can execute various deep neural networks efficiently while retaining programmability of imperative programming.

## Acknowledgements

We thank the anonymous reviewers for their insightful comments. This work was supported by Samsung Advanced Institute of Technology and the AWS Machine Learning Research Awards program.

## References

- [1] *1-billion-word-language-modeling-benchmark*. <https://github.com/ciprian-chelba/1-billion-word-language-modeling-benchmark>.
- [2] *Keras: The Python Deep Learning library*. <https://keras.io/>.
- [3] *The MNIST Database of handwritten digits*. <http://yann.lecun.com/exdb/mnist/>.
- [4] lm, 2016. <https://github.com/rafaljozefowicz/lm>.
- [5] RecursiveNN, 2016. <https://github.com/sapruash/RecursiveNN>.
- [6] ABADI, M., ET AL. TensorFlow: A system for large-scale machine learning. In *OSDI* (2016).
- [7] BEZANSON, J., EDELMAN, A., KARPINSKI, S., AND SHAH, V. B. *Julia: A Fresh Approach to Numerical Computing*, 2017.
- [8] CHEN, T., LI, M., LI, Y., LIN, M., WANG, N., WANG, M., XIAO, T., XU, B., ZHANG, C., AND ZHANG, Z. MXNet: A flexible and efficient machine learning library for heterogeneous distributed systems. In *Workshop on Machine Learning Systems in NIPS* (2015).
- [9] CULLER, D. E. Dataflow architectures. *Annual review of computer science* 1, 1 (1986), 225–253.
- [10] DEAN, J., AND GHEMAWAT, S. Mapreduce: simplified data processing on large clusters. *Communications of the ACM* 51, 1 (2008), 107–113.
- [11] DENNIS, J. B., AND MISUNAS, D. P. A preliminary architecture for a basic data-flow processor. In *ACM SIGARCH Computer Architecture News* (1975), vol. 3, ACM, pp. 126–132.
- [12] FACEBOOK. Caffe2, 2017. <https://caffe2.ai>.
- [13] FOUNDATION, P. S. *Python programming language*. <https://www.python.org/>.
- [14] FROSTIG, R., JOHNSON, M. J., AND LEARY, C. Compiling machine learning programs via high-level tracing. In *SysML* (2018).
- [15] GOYAL, P., DOLLÁR, P., GIRSHICK, R. B., NOROHDUIS, P., WESOLOWSKI, L., KYROLA, A., TULLOCH, A., JIA, Y., AND HE, K. Accurate, large minibatch SGD: training imagenet in 1 hour.
- [16] GROPP, W., LUSK, E., AND SKJELLUM, A. *Using MPI: portable parallel programming with the message-passing interface*, vol. 1. MIT press, 1999.
- [17] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition. In *CVPR* (2016).
- [18] IOFFE, S., AND SZEGEDY, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
- [19] ISARD, M., BUDIU, M., YU, Y., BIRRELL, A., AND FETTERLY, D. Dryad: distributed data-parallel programs from sequential building blocks. In *ACM SIGOPS operating systems review* (2007), vol. 41, ACM, pp. 59–72.
- [20] JEONG, E., JEONG, J. S., KIM, S., YU, G.-I., AND CHUN, B.-G. Improving the expressiveness of deep learning frameworks with recursion. In *EuroSys* (2018).
- [21] JOZEFOWICZ, R., VINYALS, O., SCHUSTER, M., SHAZEER, N., AND WU, Y. Exploring the limits of language modeling.
- [22] LECUN, Y., BUTTOU, L., BENGIO, Y., AND HAFNER, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* (1998).
- [23] LOOKS, M., HERRESHOFF, M., HUTCHINS, D., AND NORVIG, P. Deep learning with dynamic computation graphs. In *ICLR* (2017).
- [24] MXNET. *Deep Learning Programming Style*, 2018. [https://mxnet.incubator.apache.org/architecture/program\\_model.html](https://mxnet.incubator.apache.org/architecture/program_model.html).

- [25] MXNET DEVELOPERS. *Gluon*, 2018. <http://gluon.mxnet.io/>.
- [26] MÜLLER, S. C., ALONSO, G., AND CSILLAGHY, A. A. A. Pydron: Semi-automatic parallelization for multi-core and the cloud. In *OSDI* (2014).
- [27] NEUBIG, G., ET AL. DyNet: The Dynamic Neural Network Toolkit, 2017. *arXiv preprint arXiv:1701.03980*.
- [28] NIKHIL, R., ET AL. Executing a program on the mit tagged-token dataflow architecture. *IEEE Transactions on computers* 39, 3 (1990), 300–318.
- [29] PASZKE, A., GROSS, S., CHINTALA, S., CHANAN, G., YANG, E., DEVITO, Z., LIN, Z., DESMAISON, A., ANTIGA, L., AND LERER, A. Automatic differentiation in pytorch. In *Autodiff Workshop in NIPS* (2017).
- [30] PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V., ET AL. Scikit-learn: Machine learning in python. *Journal of machine learning research* 12, Oct (2011), 2825–2830.
- [31] RUSSAKOVSKY, O., DENG, J., SU, H., KRAUSE, J., SATHEESH, S., MA, S., HUANG, Z., KARPATHY, A., KHOSLA, A., BERNSTEIN, M., ET AL. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [32] SERGEEV, A., AND BALSIO, M. D. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799* (2018).
- [33] SOCHER, R., LIN, C. C.-Y., NG, A. Y., AND MANNING, C. D. Parsing natural scenes and natural language with recursive neural networks. In *ICML* (2011).
- [34] SOCHER, R., PERELYGIN, A., WU, J., CHUANG, J., MANNING, C. D., NG, A., AND POTTS, C. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP* (2013).
- [35] SZEGEDY, C., VANHOUCHE, V., IOFFE, S., SHLENS, J., AND WOJNA, Z. Rethinking the inception architecture for computer vision. In *CVPR* (2016).
- [36] TAI, K. S., SOCHER, R., AND MANNING, C. D. Improved semantic representations from tree-structured long short-term memory networks. In *ACL* (2015).
- [37] TENSORFLOW. *Eager Execution*, 2018. [https://www.tensorflow.org/programmers\\_guide/eager](https://www.tensorflow.org/programmers_guide/eager).
- [38] TENSORFLOW. *Swift for TensorFlow*, 2018. <https://github.com/tensorflow/swift>.
- [39] TENSORFLOW. *tf.contrib.eager.defun*, 2018. [https://www.tensorflow.org/versions/master/api\\_docs/python/tf/contrib/eager/defun](https://www.tensorflow.org/versions/master/api_docs/python/tf/contrib/eager/defun).
- [40] THEANO DEVELOPMENT TEAM. Theano: A Python Framework for Fast Computation of Mathematical Expressions, 2016. *arXiv preprint arXiv:1605.02688*.
- [41] WALT, S. V. D., COLBERT, S. C., AND VAROQUAUX, G. The NumPy Array: a Structure for Efficient Numerical Computation, 2011. *Computing in Science & Engineering*, 13, 2, 22-30.
- [42] YU, D., EVERSOLE, A., SELTZER, M., YAO, K., HUANG, Z., GUENTER, B., KUCHARIEV, O., ZHANG, Y., SEIDE, F., WANG, H., ET AL. An introduction to computational networks and the computational network toolkit. *Microsoft Technical Report MSR-TR-2014-112* (2014).
- [43] YU, Y., ET AL. Dynamic control flow in large-scale machine learning. In *EuroSys* (2018).
- [44] ZAREMBA, W., SUTSKEVER, I., AND VINYALS, O. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329* (2014).

## Appendix

### A Python Syntax Coverage

Table 4 describes the entire set of opcode in the CPython [13] 3.5.2 interpreter, and maps them to the sections which describe the corresponding graph generation rules. Python programs whose opcodes are mapped to Section 4.3 can only be executed on the imperative executor, and the others can be executed on the graph executor if they satisfy the assumptions. Also, Python

features that are not covered in previous sections are discussed in Appendix A.1.

## A.1 Other Dynamic Features

**Exceptions.** A Python `raise` statement can be represented as an `AssertOp` in the dataflow graph. When the `AssertOp` for an exception aborts the graph execution, the fallback occurs, and the actual, Python-style exception can be safely raised on the imperative executor. Under the same principle, for `try-except-finally` statements, only the `try-finally` part is converted into the graph elements, and the `except` part is simply not converted, since the exception will never be caught by the symbolic graph. By avoiding exception handling inside the symbolic graph, we can protect users from having to debug through symbolic graph execution traces, which are relatively more complicated than imperative execution traces.

**Context manager.** Since exception handling always occurs on the imperative executor, as described in the previous paragraph, the `with` statement can be converted into the simple function calls to `__enter__` and `__exit__` of the corresponding context manager object.

Opcode	Num	Description	Section Reference
POP_TOP, ROT_TWO, ROT_THREE, DUP_TOP, DUP_TOP_TWO, NOP, EXTENDED_ARG	7	Stack manipulation	No conversion is necessary
LOAD_CONST	1	constant	Section 4.1
UNARY_INVERT, UNARY_NEGATIVE, UNARY_NOT, UNARY_POSITIVE, BINARY_ADD, BINARY_AND, BINARY_FLOOR_DIVIDE, BINARY_LSHIFT, BINARY_MATRIX_MULTIPLY, BINARY_MODULO, BINARY_MULTIPLY, BINARY_OR, BINARY_POWER, BINARY_RSHIFT, BINARY_SUBTRACT, BINARY_TRUE_DIVIDE, BINARY_XOR, INPLACE_ADD, INPLACE_AND, INPLACE_FLOOR_DIVIDE, INPLACE_LSHIFT, INPLACE_MATRIX_MULTIPLY, INPLACE_MODULO, INPLACE_MULTIPLY, INPLACE_OR, INPLACE_POWER, INPLACE_RSHIFT, INPLACE_SUBTRACT, INPLACE_TRUE_DIVIDE, INPLACE_XOR, COMPARE_OP	31	Mathematical operators	Section 4.1
LOAD_FAST, STORE_FAST, DELETE_FAST, UNPACK_SEQUENCE, UNPACK_EX	5	local variables	Section 4.1
JUMP_ABSOLUTE, JUMP_FORWARD, JUMP_IF_FALSE_OR_POP, JUMP_IF_TRUE_OR_POP, POP_JUMP_IF_FALSE, POP_JUMP_IF_TRUE, POP_BLOCK, GET_ITER, FOR_ITER, BREAK_LOOP, CONTINUE_LOOP, SETUP_LOOP	12	dynamic control flow	Section 4.2.1
CALL_FUNCTION, CALL_FUNCTION_KW, CALL_FUNCTION_VAR, CALL_FUNCTION_VAR_KW, RETURN_VALUE, MAKE_FUNCTION	6	function call	Section 4.2.1, Section 4.3.1
LOAD_ATTR, STORE_ATTR, DELETE_ATTR	3	arbitrary object	Section 4.2.2, Section 4.2.3
BUILD_LIST, BUILD_LIST_UNPACK, LIST_APPEND, BUILD_MAP, BUILD_MAP_UNPACK, BUILD_MAP_UNPACK_WITH_CALL, MAP_ADD, BUILD_SET, BUILD_SET_UNPACK, SET_ADD, BUILD_SLICE, BUILD_TUPLE, BUILD_TUPLE_UNPACK, BINARY_SUBSCR, STORE_SUBSCR, DELETE_SUBSCR	16	list, set, map	Section 4.2.2, Section 4.2.3
LOAD_GLOBAL, LOAD_DEREF, LOAD_NAME, STORE_GLOBAL, STORE_DEREF, STORE_NAME, DELETE_GLOBAL, DELETE_DEREF, DELETE_NAME, LOAD_CLOSURE, MAKE_CLOSURE	11	non-local variables	Section 4.2.3
POP_EXCEPT, SETUP_EXCEPT, SETUP_FINALLY, RAISE_VARARGS, END_FINALLY	5	exception handling	Section A.1
SETUP_WITH, WITH_CLEANUP_FINISH, WITH_CLEANUP_START	3	with	Section A.1
YIELD_FROM, YIELD_VALUE, GET_YIELD_FROM_ITER	3	yield	Section 4.3.2
IMPORT_FROM, IMPORT_NAME, IMPORT_STAR	3	in-line import	Section 4.3.2
LOAD_BUILD_CLASS, LOAD_CLASSDEREF	2	in-line class definition	Section 4.3.2
GET_AITER, GET_ANEXT, GET_AWAITABLE, BEFORE_ASYNC_WITH, SETUP_ASYNC_WITH	5	coroutine	Section 4.3.2
Total	113		

Table 4: The mapping of the full list of CPython opcode and the corresponding sections.