

---

# A LIBRARY FOR REPRESENTING PYTHON PROGRAMS AS GRAPHS FOR MACHINE LEARNING

---

**David Bieber\***  
Google Research

**Kensen Shi**  
Google Research

**Petros Maniatis**  
Google Research

**Charles Sutton**  
Google Research

**Vincent Hellendoorn**  
Carnegie Mellon University

**Daniel Johnson**  
Google Research

**Daniel Tarlow**  
Google Research

## ABSTRACT

Graph representations of programs are commonly a central element of machine learning for code research. We introduce an open source Python library `python_graphs` that applies static analysis to construct graph representations of Python programs suitable for training machine learning models. Our library admits the construction of control-flow graphs, data-flow graphs, and composite “program graphs” that combine control-flow, data-flow, syntactic, and lexical information about a program. We present the capabilities and limitations of the library, perform a case study applying the library to millions of competitive programming submissions, and showcase the library’s utility for machine learning research.

## 1 Introduction

In this report we present `python_graphs`<sup>1</sup>, a Python library for constructing graph representations of Python programs for use in machine learning research. This report details the capabilities and limitations of the library as they pertain to applying machine learning to source code.

A standard class of approaches in applying machine learning to code is to construct a graph representation of a program, and then to perform the analysis of interest on that graph representation, learning from a large dataset of labeled example programs. Graph representations of programs used for machine learning include the abstract syntax tree (AST), control-flow graph (CFG), data-flow graphs, inter-procedural control-flow graph (ICFG), interval graph, and composite “program graphs” that encode information from multiple of the aforementioned graphs, possibly with additional program-derived data.

The `python_graphs` library directly allows for the construction of some of these graph types (e.g., control-flow graphs and composite program graphs) from arbitrary Python programs, and it provides tools that aid in constructing the others. It has been used successfully in a variety of machine learning for code publications, and we make it available as free and open source software to allow for broader use.

In Section 2 we present an overview of the use of graph representations of code in machine learning. In Section 3 we describe the capabilities (Section 3.1), possible extensions (Section 3.2), and limitations (Section 3.3) of `python_graphs`. Section 4 highlights the applications of `python_graphs` for machine learning research. Section 5 presents a case study applying `python_graphs` to 3.3 million programs from Project CodeNet [28].

## 2 Background

**Graph representations of code in machine learning** Graph representations of source code are regularly used in machine learning research. Most common among these is the abstract syntax tree. Several works learn directly from ASTs [4–6, 10, 18, 23, 25, 35, 36, 38, 41, 42, 45] or produce an AST as output [29, 40], while Johnson et al. [16] learns to dynamically augment an AST with new edges useful for a downstream task. Other works operate on a program’s

---

\*Correspondence to: [dbieber@google.com](mailto:dbieber@google.com)

<sup>1</sup><https://github.com/google-research/python-graphs>

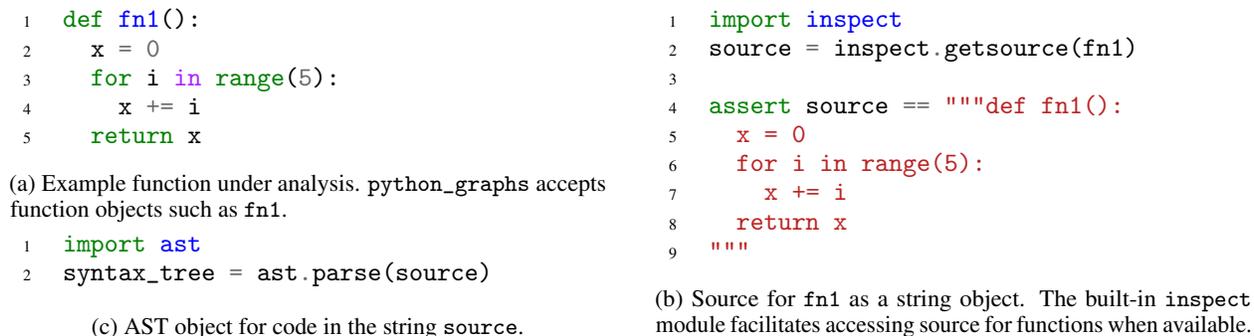


Figure 1: The input formats accepted by the `python_graphs` library are (a) function, (b) source code, and (c) AST. The code snippets here demonstrate construction of each input format for the example function `fn1`.

control-flow graph [7, 8, 11, 34] or data-flow graph [14, 17], or joint control and data flow graph (CDFG) [33]. A typical composite program graph uses an AST backbone with some subset of control-flow, data-flow, lexical, and syntactic information encoded as additional edges [2, 3, 15, 21, 22, 39, 43, 44]. Swarna et al. [30] meanwhile uses AST, CFG, and program dependence graph (PDG) representations concurrently, without unifying them into a single graph. Pashakhanloo et al. [27] forms a graph via CodeQL queries to a database representing a program. Georgiev et al. [12] forms a hypergraph, where edges can connect more than two nodes, containing again control-flow, data-flow, lexical, and syntactic information. Still other representations include a program’s interval graph [37] or a graph formed from the pointers in the heap [24]. A graph can also encode additional information, e.g., as in Tarlow et al. [31] which constructs a graph jointly representing code, a compiler error, and a build file.

Our work directly admits constructing control-flow graphs, performing data-flow analyses, and constructing certain composite program graphs from Python programs (Section 3.1). It can also be extended for constructing interprocedural control-flow graphs, novel composite program graphs, additional data-flow graphs, or span-mapped graphs (Section 3.2).

**Tools for constructing graph representations** We compare `python_graphs` with existing Python static analysis tools. Tree-sitter [9] can build a concrete syntax tree for a given source file and update it incrementally as the source changes. It supports over 40 languages including Python. Our system must operate directly on the built-in Python AST rather than a language agnostic syntax tree. CodeQL [26] is a query language for source code. These queries admit searching for control-flow and data-flow paths in source code. `pycfg` [13] generates control-flow graphs from Python source in a similar manner to `python_graphs`, but lacks support for certain language features like exceptions and generators. Scalpel [20] similarly generates control-flow graphs from Python and also performs additional static analyses, e.g., call graph construction. `python_graphs` performs data-flow analyses on top of its control-flow graphs, producing composite program graphs containing control-flow, data-flow, syntactic, and lexical information in one graph.

### 3 Capabilities, Possible Extensions, and Limitations

We provide a comprehensive overview of the capabilities of the `python_graphs` library, a discussion of how `python_graphs` can enable still further capabilities (i.e. assisting in constructing the graph types not directly supported by the library today), and a discussion of the library’s limitations.

#### 3.1 Capabilities

The `python_graphs` library enables a number of static analyses on Python source code. The main use cases are computing control-flow graphs, performing data-flow analyses, constructing composite “program graphs,” and measuring cyclomatic complexity of Python programs and functions [1]. Each of these operations may be applied to a full Python program or an individual Python function. The library handles any of the following input types: (a) Python function, (b) source code string, or (c) abstract syntax tree. Figure 1 shows constructing all three valid input formats for a sample program. In all cases, the library first converts the input to an abstract syntax tree for analysis.

##### 3.1.1 Control-Flow Graphs

A control-flow graph represents the possible paths of execution through a program. Each node in a control-flow graph represents a basic block. A basic block is a straight-line section of source code that is executed contiguously. The

```

1 # 1. Use control_flow to construct a CFG.
2 from python_graphs import control_flow
3 graph = control_flow.get_control_flow_graph(fn1)

1 # 2. Access a particular basic block by source.
2 block = graph.get_block_by_source("x += i")

1 # 3. Convert the CFG to a pygraphviz.AGraph.
2 from python_graphs import control_flow_graphviz
3 agraph = control_flow_graphviz.to_graphviz(graph)

```

(a) Example usage of `python_graphs` to (1) construct a CFG, (2) access basic blocks by source, and (3) convert to `pygraphviz` for visualization.

n	Source	Control-Flow Graph
1	<code>x = 0</code>	
2	<code>.0 = range(5)</code>	
3	<code>i = next(.0)</code>	
4	<code>x += i</code>	
5	<code>return x</code>	

(b) The statement-level control-flow graph for the program `fn1` introduced in Figure 1. `.0` denotes the iterator constructed by the for loop in `fn1`.

Figure 2: `python_graphs` supports construction and analysis of control-flow graphs for arbitrary Python functions.

only branches into a basic block enter at the start, and the only branches out of a basic block exit at the end (other than Exceptions). An edge in a control-flow graph represents a possible path of execution. There is an edge between node A and node B in a program’s control-flow graph if and only if it is possible to execute basic block B immediately following the conclusion of executing basic block A [1].

In addition to producing these standard control-flow graphs, the `python_graphs` library can also produce statement-level control-flow graphs. A node in a statement-level control-flow graph represents a single line or instruction, rather than a complete basic block. An edge between two nodes indicates that the two lines may be executed in succession. Figure 2b shows the statement-level CFG for the `fn1` program.

A control-flow graph is useful for machine learning for source code in two respects. First, it is a useful representation of code suitable for processing with graph neural networks, for example in Bieber et al. [7, 8]. Second, the control-flow graph forms the basis for a number of further analyses including data-flow analyses (liveness analysis, reachability, etc.), computing cyclomatic complexity, and constructing program graphs, each implemented by the `python_graphs` library.

In Python, any line of code can raise an exception. Taking this form of execution into account, this limits basic blocks to a single line of code, since a raised exception is an exit branch. Rather than restrict basic blocks to a single line of code, we take a more pragmatic approach, and introduce a second optional edge type, “interrupting edges”, in our control-flow graph data structure that represents control flow due to exceptions. An interrupting edge from block A to block B indicates that an exception raised during the execution of A can cause control to flow to block B. `python_graphs` control-flow graphs can be used with or without these interrupting edges.

To construct a control-flow graph with `python_graphs`, use the `control_flow` module’s `get_control_flow_graph` as in Figure 2.

### 3.1.2 Data-Flow Analyses

A data-flow analysis computes information about how the variables in the program are used, such as which variables are *live* at a given program location. A live variable is one that may be read at some point in the future before its value is overwritten. The `python_graphs` library implements two best-effort data-flow analyses: liveness and last-access analysis.

Data-flow analyses are performed through iterative application of the data-flow equations until a fixed point is reached. The `python_graphs` library supports both forward and backward data-flow analysis, and so can be extended to support additional data-flow analyses. Liveness is implemented as a backward analysis, and last-access as a forward analysis.

An example of using the liveness analysis to obtain the set of loop variables in a while loop is provided with the library, a necessary step in rewriting Python while loops into their functional form. The `data_flow` module provides the data-flow analyses, as in Figure 3.

### 3.1.3 Composite Program Graphs

The `python_graphs` library implements a single kind of composite program graph, based closely on that of Allamanis et al. [2]. In this document we refer to these composite program graphs simply as “program graphs”, though of course other kinds of program graphs are possible, with different node and edge types.

```

1 # 1. Use data_flow to perform liveness analysis.
2 from python_graphs import data_flow
3 analysis = data_flow.LivenessAnalysis()
4 for block in graph.get_exit_blocks():
5     analysis.visit(block)

1 # 2. Construct a program graph with program_graph.
2 from python_graphs import program_graph
3 pg = program_graph.get_program_graph(program)

1 # 3. Access a particular node by source.
2 node = pg.get_node_by_source_and_identifier(
3     'return x', 'x')
```

(a) Example usage of `python_graphs` to (1) perform liveness analysis on a program’s control-flow graph. Independently, (2) shows constructing a composite program graph, and (3) accessing one of its node by source.

#	Source	Live in	Live out
1	<code>x = 0</code>	$\emptyset$	
	<code>.0 = range(5)</code>		$\{x\}$
2	<code>i = next(.0)</code>	$\{x\}$	$\{x, i\}$
3	<code>x += i</code>	$\{x, i\}$	$\{x\}$
4	<code>return x</code>	$\{x\}$	$\emptyset$

(b) The results of the liveness data-flow analysis. Live in and live out indicate the variables that are live at the start and end of each basic block respectively. # denotes basic block number.

Figure 3: Example usage of data-flow analysis and program graph construction in `python_graphs`.

A program graph has the abstract syntax tree of the program it represents as its backbone. Each node in the program graph directly corresponds to a single node in the AST, and vice versa. Lists and primitive values in the AST have corresponding nodes in the program graph as well. Corresponding to each syntax element in the program (leaf nodes in the AST) is a syntax node in the program graph. Each edge in the AST also appears in the program graph. The program graph then has additional edges representing the following relationships between program pieces: “NEXT\_SYNTAX”, “LAST\_LEXICAL\_USE”, “CFG\_NEXT”, “LAST\_READ”, “LAST\_WRITE”, “COMPUTED\_FROM”, “CALLS”, “FORMAL\_ARG\_NAME”, and “RETURNS\_TO”. Collectively, the edges in a program graph convey control-flow, data-flow, lexical, and syntactic information about the program.

We summarize the edge types and their meanings in Table 1. These edge types are also useful for constructing other graph types (Section 3.2): interprocedural control-flow graphs, data-flow graphs, and other composite program graphs.

The `python_graphs` library provides the function `get_program_graph` for constructing a program graph from any of the supported input types (source code, an abstract syntax tree, or Python function). Figure 3a shows example usage.

Table 2 lists several programs along with their control-flow graphs and program graphs as computed by `python_graphs`.

### 3.1.4 Cyclomatic Complexity

Cyclomatic complexity is a standard measure of program complexity based on the set of possible paths through a program. It measures the number of linearly independent execution paths through a program. The `python_graphs` library can compute the cyclomatic complexity of a Python function. This functionality is available via the function `cyclomatic_complexity`, which accepts a Python function (as source, AST, or Python function object) and returns an integer. To compute the cyclomatic complexity of a program, `python_graphs` first constructs its control-flow graph. In a control-flow graph with  $E$  edges,  $N$  nodes, and  $P$  distinct connected components, the cyclomatic complexity  $M$  is given by  $M = E - N + 2P$ .

EDGE TYPE	DESCRIPTION
FIELD	<i>dest</i> is a field of AST node <i>src</i> .
NEXT_SYNTAX	<i>dest</i> is the syntax element immediately following <i>src</i> in top-to-bottom left-to-right order.
LAST_LEXICAL_USE	The variable at node <i>dest</i> has its previous appearance at <i>src</i> in top-to-bottom left-to-right order.
CFG_NEXT	The statement indicated by <i>dest</i> can be executed immediately following that indicated by <i>src</i> .
LAST_READ	When <i>src</i> is about to execute, it may be that the variable at <i>src</i> was most recently read at <i>dest</i> .
LAST_WRITE	When <i>src</i> is about to execute, it may be that the variable at <i>src</i> was most recently written to at <i>dest</i> .
COMPUTED_FROM	<i>src</i> indicates a variable on an assignment’s left hand side, and <i>dest</i> a variable on its right hand side.
CALLS	<i>src</i> is an AST call node, and <i>dest</i> is the definition of the function being called.
FORMAL_ARG_NAME	<i>src</i> is an argument in a function call; <i>dest</i> is the corresponding parameter in the function definition.
RETURNS_TO	<i>src</i> is the return node in a function definition, and <i>dest</i> is the AST call node that calls that function.

Table 1: An edge of the given edge type from node *src* to node *dest* has the meaning given in this table.

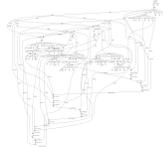
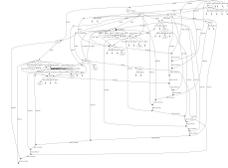
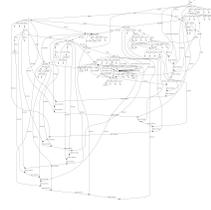
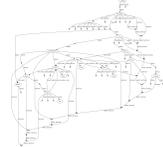
#	Source	n	Statement	CFG	Program Graph
1	<pre>def fn1(a, b):     return a + b</pre>	1 2 3	<pre>a, b ← args return a + b &lt;exit&gt;</pre>		 Figure 6
2	<pre>def fn2(a, b):     c = a     if a &gt; b:         c -= b     return c</pre>	1 2 3 4 5 6	<pre>a, b ← args c = a a &gt; b c -= b return c &lt;exit&gt;</pre>		 Figure 7
3	<pre>def fn3(a, b):     c = a     if a &gt; b:         c -= b         c += 1         c += 2         c += 3     else:         c += b     return c</pre>	1 2 3 4 5 6 7 8 9 10	<pre>a, b ← args c = a a &gt; b c -= b c += 1 c += 2 c += 3 c += b return c &lt;exit&gt;</pre>		 Figure 8
4	<pre>def fn4(i):     count = 0     for i in range(i):         count += 1     return count</pre>	1 2 3 4 5 6 7	<pre>i ← args count = 0 range(i) i ← iter count += 1 return count &lt;exit&gt;</pre>		 Figure 9
5	<pre>def fn5(i):     count = 0     for _ in range(i):         if count &gt; 5:             break         count += 1     return count</pre>	1 2 3 4 5 6 7 8	<pre>i ← args count = 0 range(i) _ ← iter count &gt; 5 count += 1 return count &lt;exit&gt;</pre>		 Figure 10
6	<pre>def fn6():     count = 0     while count &lt; 10:         count += 1     return count</pre>	1 2 3 4 5	<pre>count = 0 count &lt; 10 count += 1 return count &lt;exit&gt;</pre>		 Figure 11
7	<pre>def fn7():     try:         raise ValueError('N/A')     except ValueError as e:         del e     return</pre>	1 2 3 4 5 6	<pre>raise ValueError('N/A') ValueError e ← exception del e return &lt;exit&gt;</pre>		 Figure 12
8	<pre>def fn8(a):     a += 1</pre>	1 2 3	<pre>a ← args a += 1 &lt;exit&gt;</pre>		 Figure 13

Table 2: Example programs and their associated control-flow graphs and program graphs. Enlarged versions of the program graph figures are included in Appendix A.

## 3.2 Possible Extensions

The following capabilities are possible to implement using `python_graphs`, but are not directly provided by `python_graphs` out of the box.

### 3.2.1 Alternative Composite Program Graphs

The program graphs generated by `python_graphs`'s `get_program_graph` make specific choices for what nodes and edges are included in the graph. Other choices are possible. For alternative composite program graph schemes, the source of `python_graphs`'s `get_program_graph` function serves as an illustrative example for how to construct a composite program graph with the desired set of nodes and edges.

### 3.2.2 Inter-procedural Control-Flow Graphs

`python_graphs`'s `get_control_flow_graph` function constructs the control-flow graph for a single function or program; it does not include edges indicating control flow between functions.

An interprocedural control-flow graph (ICFG) is a control-flow graph that shows the control flow possible between functions, not just within a function. We can view an ICFG as a composite program graph consisting of the control-flow graphs of a program and all its functions, as well as `CALLS` and `RETURNS_TO` edges indicating the possible interprocedural control flows in the program. `python_graphs` provides the capability for constructing the necessary control-flow graphs and additional edges, making it possible to write a function to construct ICFGs as well.

### 3.2.3 Data-Flow Graphs

A data-flow graph represents the data dependencies present in a program. The nodes in a data-flow graph represent the variable access locations in a program, and the edges in a data-flow graph denote relationships between these accesses. An example of such a relationship is an edge with `dest` indicating where a variable is assigned to and `src` indicating where the assigned value is subsequently used (equivalent to `python_graphs`'s `LAST_WRITE` edges). We can therefore view data-flow graphs as composite program graphs consisting of a subset of AST nodes (just those representing variable accesses) and selected edge types like `LAST_READ` or `LAST_WRITE`.

### 3.2.4 Span-Mapped Graphs

In order to use the graphs produced by the `python_graphs` library for machine learning applications, it can be useful to tokenize the sections of code corresponding to each node. We suggest two approaches to handling this: (1) whole program tokenization and (2) per-node tokenization.

In *whole program tokenization* we tokenize the entire program first. Then, using `python_graphs`, we create a graph structure for the program. Finally we extract for each node the span of tokens from the whole program tokenization corresponding to that node. This approach allows for the possibility that a token consisting of multiple characters will be part of two consecutive nodes, and we must choose which node(s) to associate that token with.

In *per-node tokenization*, we instead split the program source into chunks according to which node in the graph they are part of, and then tokenize those chunks independently.

The key data required by these approaches is a mapping from a graph node to a span in the textual representation of program source (approaches 1 and 2) or to a span in the tokenized representation of program source (approach 1 only). We call a graph augmented with this data a *span-mapped graph*. Both approaches are possible using `python_graphs`. Bieber et al. [8] uses approach 1, with code freely available online<sup>2</sup>. This same code example is informative for any project wishing to implement approach 2.

### 3.2.5 Additional Data-Flow Analyses

`python_graphs` implements liveness and last-access data-flow analyses, and provides a framework for implementing additional analyses. This framework allows somewhat straightforward implementation of definite assignment analysis or computing reaching definitions, for example.

<sup>2</sup><https://github.com/google-research/runtime-error-prediction>

### 3.3 Limitations

Python source code is difficult to analyze statically because so much of Python’s behavior is determined dynamically. We perform a “best-effort” analysis, which we do not guarantee will be correct considering all of Python’s language features. Inspection in Python allows manipulation of stack frames or of local or global variables, causing hard-to-detect effects on data and control flow. Eval and exec permit the execution of arbitrary code constructed dynamically and inaccessible to our analysis. Operations can be overloaded dynamically, so e.g. even a simple addition operation can have effects overlooked by our analyses. These language features are empowering to Python users, but restrict the guarantees our analyses can provide.

## 4 Use cases

We next show how the `python_graphs` library is used in machine learning research. Uses include both building graph representations of programs as inputs to neural networks, and providing supervision for models that output graphs.

### 4.1 Graph Representations as Model Inputs

**Instruction Pointer Attention Graph Neural Networks** The instruction pointer attention graph neural network (IPA-GNN) model family [7, 8] operates on control-flow graphs as its primary input. IPA-GNN architectures then perform a soft execution of the input program, maintaining a soft instruction pointer representing at each model step a probability distribution over the statements in a program. These works use `python_graphs` to produce the control-flow graphs for the programs under consideration, which include both simple synthetic programs [7] and complex human-authored programs from a competition setting [8].

The original IPA-GNN work [7] uses the control-flow edges as produced by `python_graphs`’s default settings, and represents each statement with a 4-tuple of values, which is possible because the domain of statements is restricted. By contrast, the follow-up work on competition programs [8] uses a larger control-flow graph that additionally includes interrupting edges, indicating to where control would flow from each node if an exception were raised during the execution of that node. Further, a sequence of tokens is associated with each node in the control-flow graph, following Section 3.2.4, allowing it to handle arbitrary human-authored Python statements.

**Global Relational Models of Source Code** Hellendoorn et al. [15] investigates models that combine global information (like a Transformer) and structural information (like a GNN), i.e., Graph-Sandwich models and the GREAT (Graph-Relational Embedding Attention Transformer) model. This paper uses `python_graphs` to construct the composite program graphs of Section 3.1.3. The models accept these program graphs as input and uses them to identify variable misuse bugs.

### 4.2 Program Graphs as Targets

**Graph Finite-State Automaton (GFSA) Layers** Johnson et al. [16] introduces a neural network layer that adds a new learned edge type to an input graph. Toward learning static analyses of Python code, it trains a neural model to take an AST as input and predict a composite program graph as output. The model thereby learns to perform both control-flow and data-flow analyses from data. For its targets, it produces a composite program graph with `python_graphs`, selecting a subset of the default edge types and introducing a few additional edge types (as in Section 3.2.1).

**Learning to Extend Program Graphs to Work-in-Progress Code** Li et al. [22] learns to predict edge relations from work-in-progress code, even when the code does not parse. The composite program graphs of Section 3.1.3 form the ground truth edge relation targets for this work.

## 5 Case Study: Project CodeNet

In order to evaluate the `python_graphs` library on the diversity of language features found in realistic code, we obtain a dataset of 3.3 million programs from Project CodeNet [28]. For each program, we use `python_graphs` to construct a control-flow graph and a composite program graph complete with syntactic, control-flow, data-flow, and lexical information about the program. We collect metrics about the resulting graphs to provide information about the robustness of `python_graphs` and the size, complexity, and connectedness of the program graphs it produces.

`python_graphs` cannot construct graph representations for every submission in Project CodeNet, as many of them do not parse. Table 3 shows how many of the programs graph construction succeeds for, as well as the failure reasons for the

Status	# Programs	Freq. (%)
Success	3,157,463	96.08
ast.parse failures	126,751	3.86
SyntaxError	114,817	3.49
IndentationError	8,893	0.27
TabError	3,032	0.09
RecursionError	5	0.00
ValueError	4	0.00
RuntimeError	2,100	0.06
return outside function	1,719	0.05
break outside loop	330	0.01
continue outside loop	51	0.00
<b>Total</b>	<b>3,286,314</b>	<b>100%</b>

Table 3: Control-flow graph construction success rates on a dataset of both valid and invalid Python submissions to competitive programming problems.

Edge Type	# / Program	Freq. (%)
FIELD	370.4	100.00
SYNTAX	163.4	99.99
NEXT_SYNTAX	162.4	99.99
LAST_LEXICAL_USE	26.1	99.12
CFG_NEXT	18.1	99.06
LAST_READ	38.0	92.29
LAST_WRITE	30.6	98.83
COMPUTED_FROM	11.7	98.55
CALLS	0.5	21.37
FORMAL_ARG_NAME	0.6	12.99
RETURNS_TO	0.7	15.56

Table 4: Frequencies of edge types in the composite program graphs for the Project CodeNet Python submissions. # / Program is the average number of occurrences of the edge type across all programs. Freq. (%) shows the percent of programs exhibiting the edge type at least once.

remaining graphs. The programs for which `python_graphs` cannot produce graph representations are predominantly those which fail to parse under Python’s own parser: `ast.parse`. The majority of such programs cause the parser to raise a `SyntaxError`, `IndentationError`, or `TabError`, with just nine leading the built in parser to raise a `RecursionError` or `ValueError`. The `python_graphs` library rejects an additional 2100 programs (0.07%) because they contain either `return` outside of a function body, or `break` or `continue` outside of a loop. In total, this result gives us confidence there are no language feature corner cases that elude the `python_graphs` library and cause failures for well-formed programs that otherwise can be run by a standard Python interpreter. In Table 4 we report for each program graph edge type the fraction of programs it appears in as well as the mean number of appearances across all programs.

We next use `python_graphs` to measure the cyclomatic complexity of each of the submissions. Figure 4 plots the relationship between program length and cyclomatic complexity. We measure program length in non-empty lines of

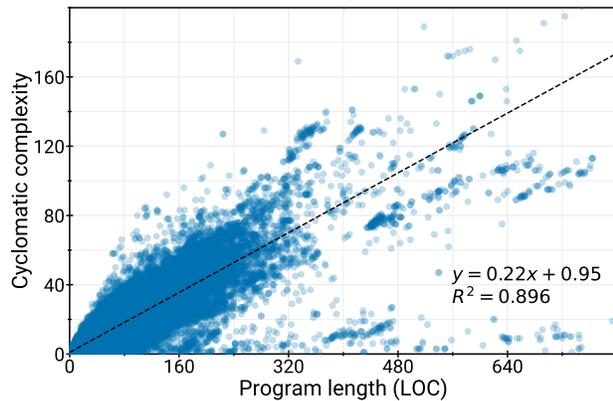


Figure 4: The relationship between program length and cyclomatic complexity for Python submissions in Project CodeNet.

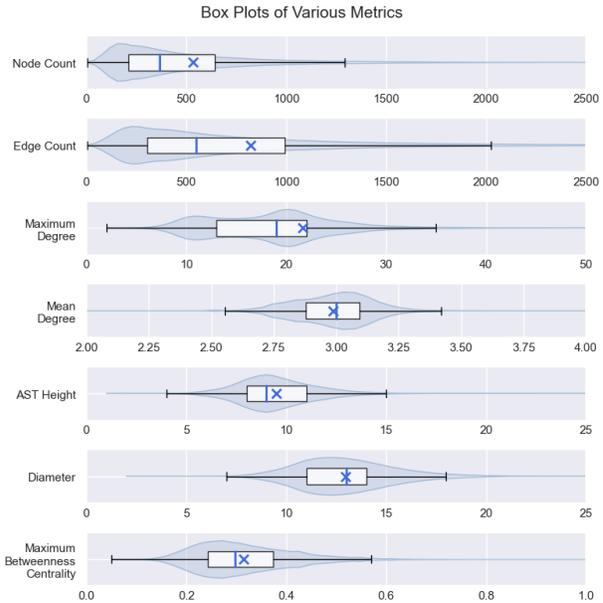


Figure 5: Box plots for various metrics of program graphs for Python submissions in Project CodeNet. The vertical blue line in each boxplot shows the median of the data as usual, while the blue  $\times$  shows the mean.

Metric	Min	Median	Mean	Max
Node Count	3	364	534.8	751,817
Edge Count	2	548	822.4	4,675,600
Maximum Degree	2	19	21.6	150,004
Mean Degree	1.3	3.0	3.0	79.9
AST Height	1	9	9.5	269
Diameter	2	13	13.0	143
Maximum Betweenness Centrality	0.0	0.3	0.3	1.0

Table 5: Summary statistics for various graph metrics across the dataset. The diameter and maximum betweenness centrality metrics do not include graphs exceeding 5000 nodes.

code (LOC). Omitting as outliers those programs longer than 800 LOC or with complexity exceeding 200 (just 118 programs out of 3.16 million), we perform linear regression and observe  $R^2 = 0.896$ , in line with prior work [19, 32].

We measure the size of program graphs according to their node counts and edge counts, the height of their AST backbone, and graph diameter. As measures of connectedness, we compute the maximum degree of a node, mean degree of the nodes, and maximum betweenness centrality of a node in the graph. The distributions of each of these metrics are shown with boxplots in Figure 5, and key summary statistics are listed in Table 5. Appendix B contains histograms showing the distribution of each metric across the dataset. These metrics convey the scale and diversity of submissions to online programming contests and the graph sizes needed for processing these submissions as python graphs with graph neural networks.

## 6 Discussion

The core capabilities of `python_graphs` for machine learning research are generating control-flow graphs, performing data-flow analyses, generating composite program graphs, and computing cyclomatic complexity of Python programs. For our research, we have been fruitfully using `python_graphs` for graph representations of programs for multiple years. The library is robust and flexible, having been successfully run on millions of programs and used in several published papers. Still, several open challenges remain for providing insights into program semantics to machine learners. First, due to the dynamic nature of Python the library’s analyses are limited to providing best-effort results, not considering the possible effects of e.g. dynamic execution or introspection. A further key limitation of the library is its restriction to processing Python programs. This makes getting a consistent graph representation across programming languages challenging, which is important when training a multi-lingual model of code. While significant recent progress has been made in machine learning for code research, many fundamental problems in the space remain open research challenges. Examples of these challenges include learning about program semantics from end-to-end program behavior, and identifying neural models exhibiting systematic generalization. For these challenges, where the structure and semantics of programs are important, `python_graphs` provides a framework to study how graph representations of programs may contribute to forward progress.

## References

- [1] Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools (2nd Edition)*. Addison-Wesley Longman Publishing Co., Inc., USA, 2006. ISBN 0321486811.
- [2] Miltiadis Allamanis, Marc Brockschmidt, and Mahmoud Khademi. Learning to represent programs with graphs. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BJOFETxR->.
- [3] Miltiadis Allamanis, Henry Jackson-Flux, and Marc Brockschmidt. Self-supervised bug detection and repair, 2021. URL <https://arxiv.org/abs/2105.12787>.
- [4] Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. code2seq: Generating sequences from structured representations of code, 2018. URL <https://arxiv.org/abs/1808.01400>.
- [5] Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. code2vec: Learning distributed representations of code, 2018. URL <https://arxiv.org/abs/1803.09473>.
- [6] Uri Alon, Roy Sadaka, Omer Levy, and Eran Yahav. Structural language models of code, 2019. URL <https://arxiv.org/abs/1910.00577>.
- [7] David Bieber, Charles Sutton, Hugo Larochelle, and Daniel Tarlow. Learning to execute programs with instruction pointer attention graph neural networks, 2020.
- [8] David Bieber, Rishab Goel, Daniel Zheng, Hugo Larochelle, and Daniel Tarlow. Static prediction of runtime errors by learning to execute programs with external resource descriptions, 2022. URL <https://arxiv.org/abs/2203.03771>.
- [9] Max Brunsfeld, Patrick Thomson, Andrew Hlynskyi, Josh Vera, Phil Turnbull, Timothy Clem, Douglas Creager, Andrew Helwer, Rob Rix, Hendrik van Antwerpen, Michael Davis, Ika, Tuan-Anh Nguyen, Stafford Brunk, Niranjan Hasabnis, bfredl, Mingkai Dong, Vladimir Pantelev, ikrima, Steven Kalt, Kolja Lampe, Alex Pinkus, Mark Schmitz, Matthew Krupcale, narpfel, Santos Gallegos, Vicent Martí, Edgar, and George Fraser. tree-sitter/tree-sitter: v0.20.6, March 2022. URL <https://doi.org/10.5281/zenodo.6326492>.
- [10] Nghi D. Q. Bui, Yijun Yu, and Lingxiao Jiang. Infercode: Self-supervised learning of code representations by predicting subtrees, 2020. URL <https://arxiv.org/abs/2012.07023>.
- [11] Daniel DeFreez, Aditya V. Thakur, and Cindy Rubio-González. Path-based function embedding and its application to specification mining, 2018. URL <https://arxiv.org/abs/1802.07779>.
- [12] Dobrik Georgiev, Marc Brockschmidt, and Miltiadis Allamanis. HEAT: Hyperedge attention networks, 2022. URL <https://arxiv.org/abs/2201.12113>.
- [13] Rahul Gopinath. pycfg: The Python control flow graph. URL <https://rahul.gopinath.org/post/2019/12/08/python-controlflow/>.
- [14] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. GraphCodeBERT: Pre-training code representations with data flow, 2020. URL <https://arxiv.org/abs/2009.08366>.
- [15] Vincent J. Hellendoorn, Charles Sutton, Rishabh Singh, Petros Maniatis, and David Bieber. Global relational models of source code. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=B1lnbRNtwr>.
- [16] Daniel D. Johnson, Hugo Larochelle, and Daniel Tarlow. Learning graph structure with a finite-state automaton layer, 2020.
- [17] Samuel J. Kaufman, Phitchaya Mangpo Phothilimthana, Yanqi Zhou, Charith Mendis, Sudip Roy, Amit Sabne, and Mike Burrows. A learned performance model for tensor processing units, 2020. URL <https://arxiv.org/abs/2008.01040>.
- [18] Seohyun Kim, Jinman Zhao, Yuchi Tian, and Satish Chandra. Code prediction by feeding trees to transformers, 2020. URL <https://arxiv.org/abs/2003.13848>.
- [19] Davy Landman, Alexander Serebrenik, and Jurgen Vinju. Empirical analysis of the relationship between CC and SLOC in a large corpus of Java methods. In *IEEE International Conference on Software Maintenance and Evolution*, pages 221–230, 2014. doi: 10.1109/ICSME.2014.44.
- [20] Li Li, Jiawei Wang, and Haowei Quan. Scalpel: The Python static analysis framework, 2022. URL <https://arxiv.org/abs/2202.11840>.

- [21] Mingzhe Li, Jianrui Pei, Jin He, Kevin Song, Frank Che, Yongfeng Huang, and Chitai Wang. Using GGNN to recommend log statement level, 2019. URL <https://arxiv.org/abs/1912.05097>.
- [22] Xuechen Li, Chris J. Maddison, and Daniel Tarlow. Learning to extend program graphs to work-in-progress code, 2021. URL <https://arxiv.org/abs/2105.14038>.
- [23] Yi Li, Shaohua Wang, and Tien N. Nguyen. Fault localization with code coverage representation learning, 2021. URL <https://arxiv.org/abs/2103.00270>.
- [24] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks, 2015. URL <https://arxiv.org/abs/1511.05493>.
- [25] Shangqing Liu, Cuiyun Gao, Sen Chen, Lun Yiu Nie, and Yang Liu. Atom: Commit message generation based on abstract syntax tree and hybrid ranking, 2019. URL <https://arxiv.org/abs/1912.02972>.
- [26] Oege de Moor, Mathieu Verbaere, Elnar Hajiyev, Pavel Avgustinov, Torbjorn Ekman, Neil Ongkingco, Damien Sereni, and Julian Tibble. Keynote address: .ql for source code analysis. In *Seventh IEEE International Working Conference on Source Code Analysis and Manipulation (SCAM 2007)*, pages 3–16, 2007. doi: 10.1109/SCAM.2007.31.
- [27] Pardis Pashakhanloo, Aaditya Naik, Yuepeng Wang, Hanjun Dai, Petros Maniatis, and Mayur Naik. CodeTrek: Flexible modeling of code using an extensible relational representation. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=WQc075jmBmf>.
- [28] Ruchir Puri, David S. Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir Choudhury, Lindsey Decker, Veronika Thost, Luca Buratti, Saurabh Pujar, Shyam Ramji, Ulrich Finkler, Susan Malaika, and Frederick Reiss. CodeNet: A large-scale AI for code dataset for learning a diversity of coding tasks, 2021. URL <https://arxiv.org/abs/2105.12655>.
- [29] Maxim Rabinovich, Mitchell Stern, and Dan Klein. Abstract syntax networks for code generation and semantic parsing, 2017. URL <https://arxiv.org/abs/1704.07535>.
- [30] Karthik Chandra Swarna, Noble Saji Mathews, Dheeraj Vagavolu, and Sridhar Chimalakonda. A mocktail of source code representations, 2021. URL <https://arxiv.org/abs/2106.10918>.
- [31] Daniel Tarlow, Subhodeep Moitra, Andrew Rice, Zimin Chen, Pierre-Antoine Manzagol, Charles Sutton, and Edward Aftandilian. Learning to fix build errors with graph2diff neural networks, 2019. URL <https://arxiv.org/abs/1911.01205>.
- [32] Yahya Tashtoush, Mohammed Al-Maolegi, and Bassam Arkok. The correlation among software complexity metrics with case study. 2014. doi: 10.48550/ARXIV.1408.4523. URL <https://arxiv.org/abs/1408.4523>.
- [33] Shobha Vasudevan, Wenjie (Joe) Jiang, David Bieber, Rishabh Singh, hamid shojaei, C. Richard Ho, and Charles Sutton. Learning semantic representations to verify hardware designs. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 23491–23504. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/file/c5aa65949d20f6b20e1a922c13d974e7-Paper.pdf>.
- [34] Anh Viet Phan, Minh Le Nguyen, and Lam Thu Bui. Convolutional neural networks over control flow graphs for software defect prediction. In *2017 IEEE 29th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 45–52, 2017. doi: 10.1109/ICTAI.2017.00019.
- [35] Wenhan Wang, Ge Li, Sijie Shen, Xin Xia, and Zhi Jin. Modular tree network for source code representation learning, 2021. URL <https://arxiv.org/abs/2104.00196>.
- [36] Yanlin Wang and Hui Li. Code completion by modeling flattened abstract syntax trees as graphs, 2021. URL <https://arxiv.org/abs/2103.09499>.
- [37] Yu Wang, Fengjuan Gao, Linzhang Wang, and Ke Wang. Learning semantic program embeddings with graph interval neural network, 2020. URL <https://arxiv.org/abs/2005.09997>.
- [38] Martin White, Michele Tufano, Christopher Vendome, and Denys Poshyvanyk. Deep learning code fragments for code clone detection. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 87–98, 2016.
- [39] Fabian Yamaguchi, Nico Golde, Daniel Arp, and Konrad Rieck. Modeling and discovering vulnerabilities with code property graphs. In *2014 IEEE Symposium on Security and Privacy*, pages 590–604, 2014. doi: 10.1109/SP.2014.44.
- [40] Pengcheng Yin and Graham Neubig. Tranx: A transition-based neural abstract syntax parser for semantic parsing and code generation, 2018. URL <https://arxiv.org/abs/1810.02720>.

- [41] Hao Yu, Wing Lam, Long Chen, Ge Li, Tao Xie, and Qianxiang Wang. Neural detection of semantic code clones via tree-based convolution. In *2019 IEEE/ACM 27th International Conference on Program Comprehension (ICPC)*, pages 70–80, 2019. doi: 10.1109/ICPC.2019.00021.
- [42] Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, Kaixuan Wang, and Xudong Liu. A novel neural source code representation based on abstract syntax tree. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, pages 783–794, 2019. doi: 10.1109/ICSE.2019.00086.
- [43] Kechi Zhang, Wenhan Wang, Huangzhao Zhang, Ge Li, and Zhi Jin. Learning to represent programs with heterogeneous graphs, 2020. URL <https://arxiv.org/abs/2012.04188>.
- [44] Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks, 2019. URL <https://arxiv.org/abs/1909.03496>.
- [45] Daniel Zügner, Tobias Kirschstein, Michele Catasta, Jure Leskovec, and Stephan Günnemann. Language-agnostic representation learning of source code from structure and context, 2021. URL <https://arxiv.org/abs/2103.11318>.

## A Program Graph Visualizations

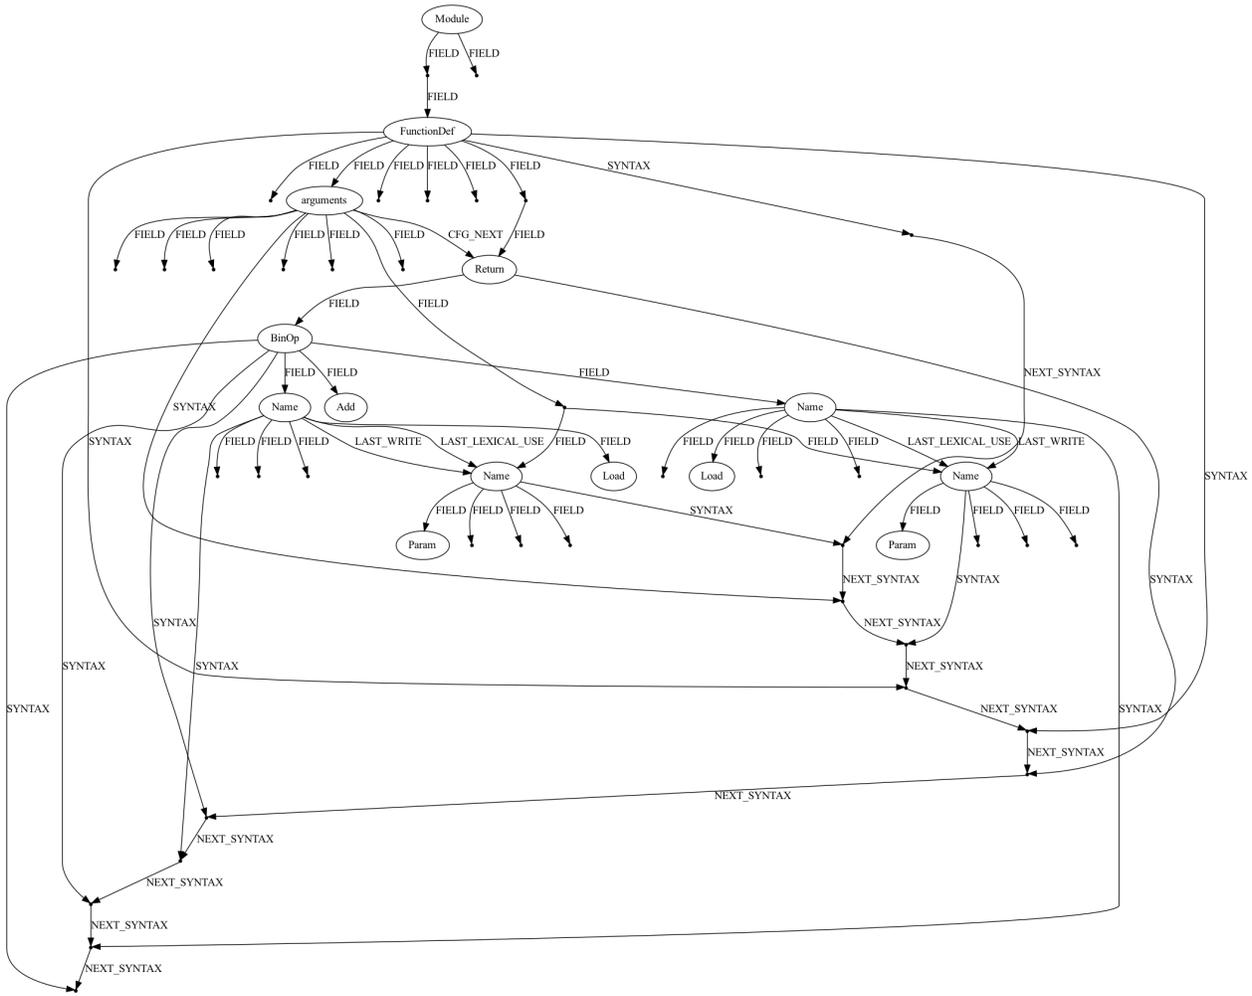


Figure 6: Program graph for Program #1 from Table 2.



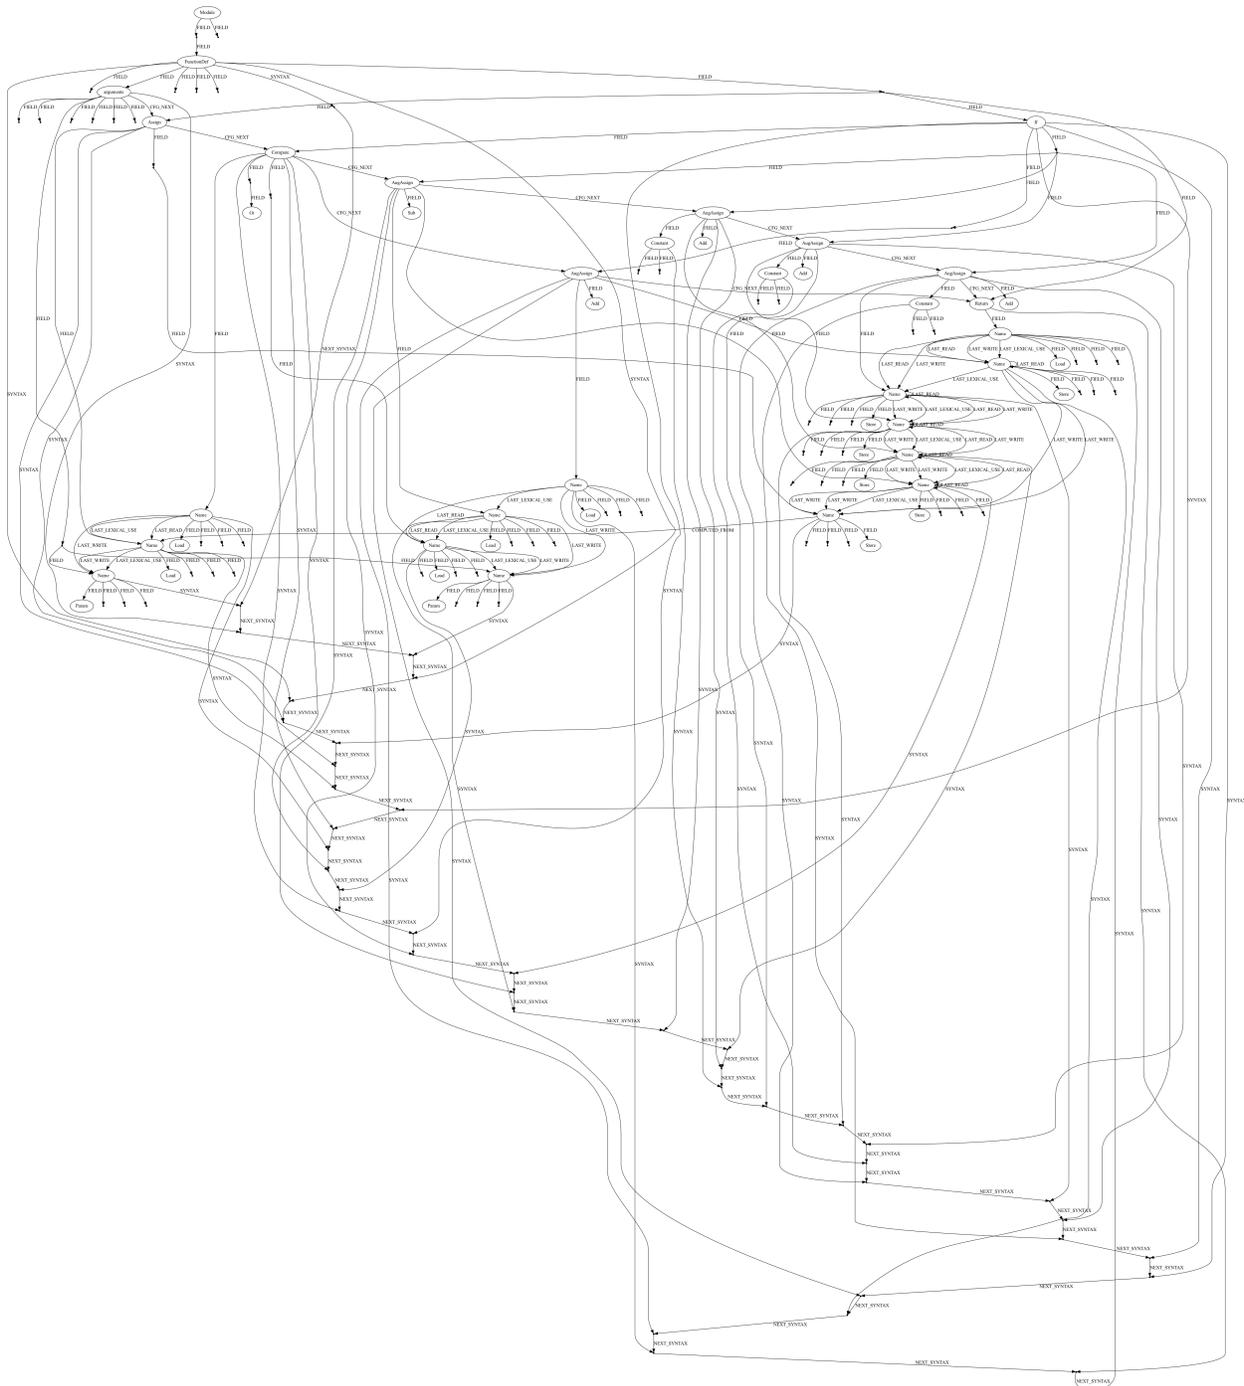


Figure 8: Program graph for Program #3 from Table 2.

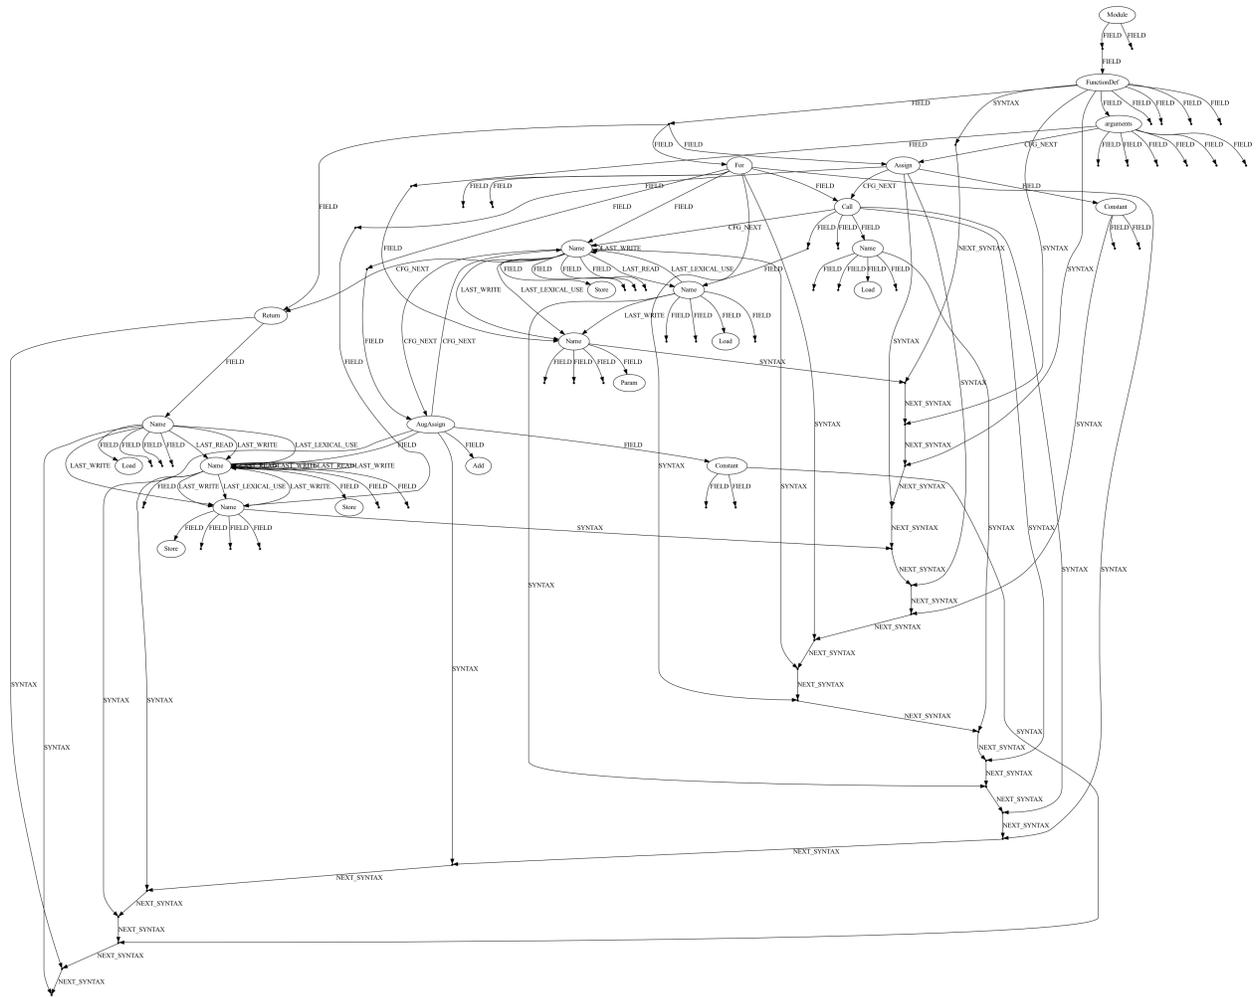


Figure 9: Program graph for Program #4 from Table 2.





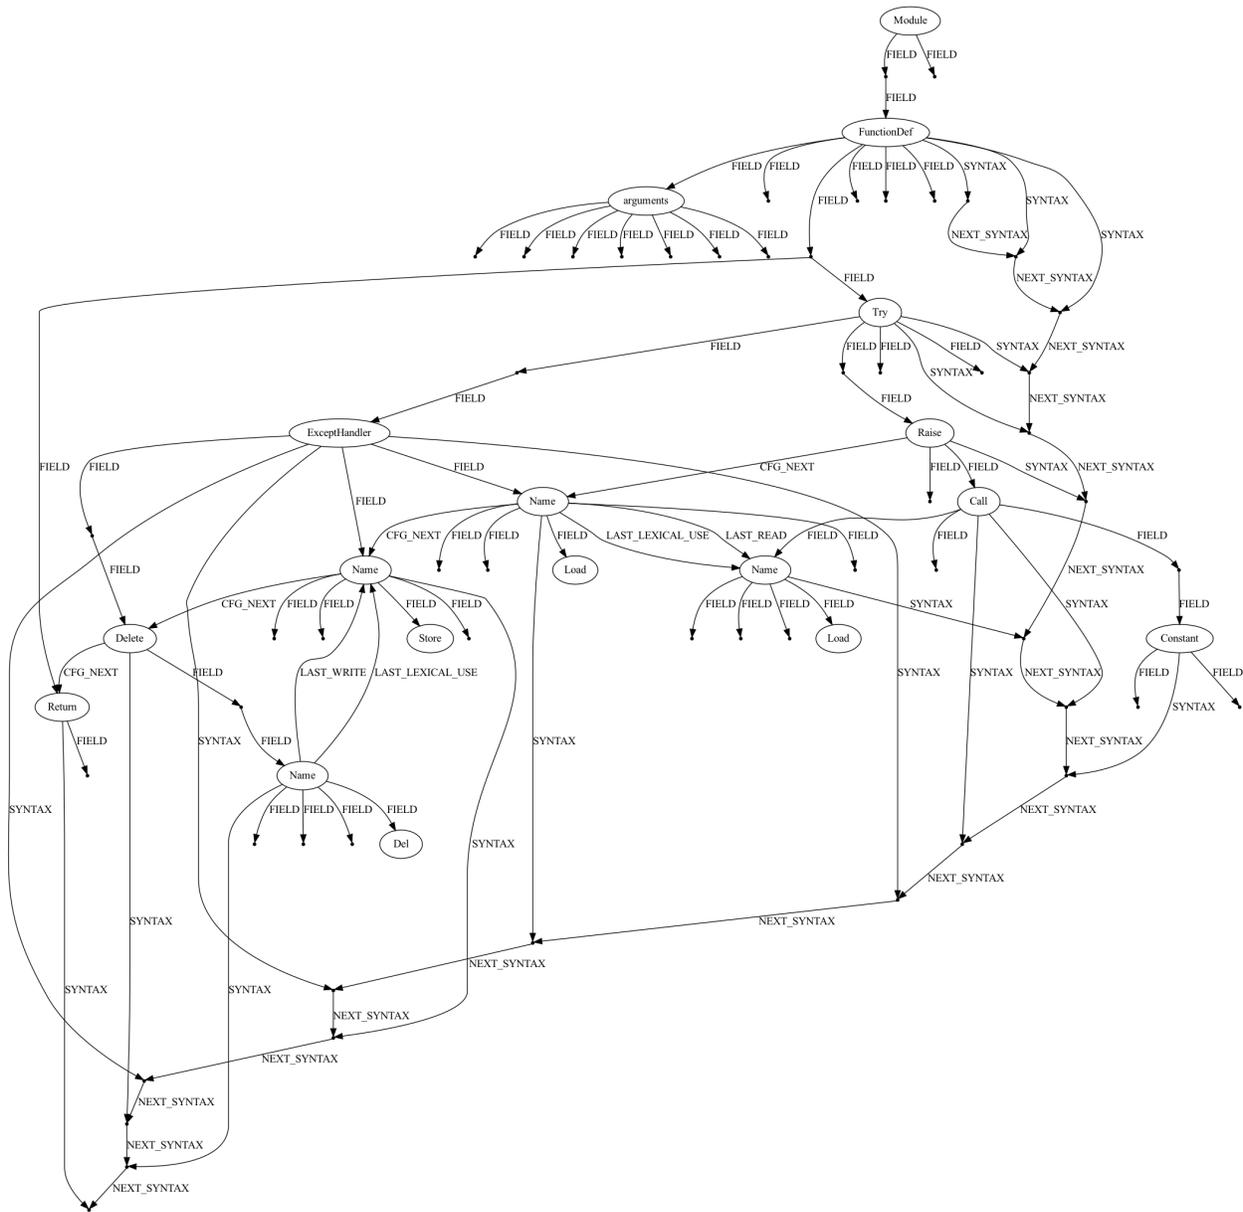


Figure 12: Program graph for Program #7 from Table 2.

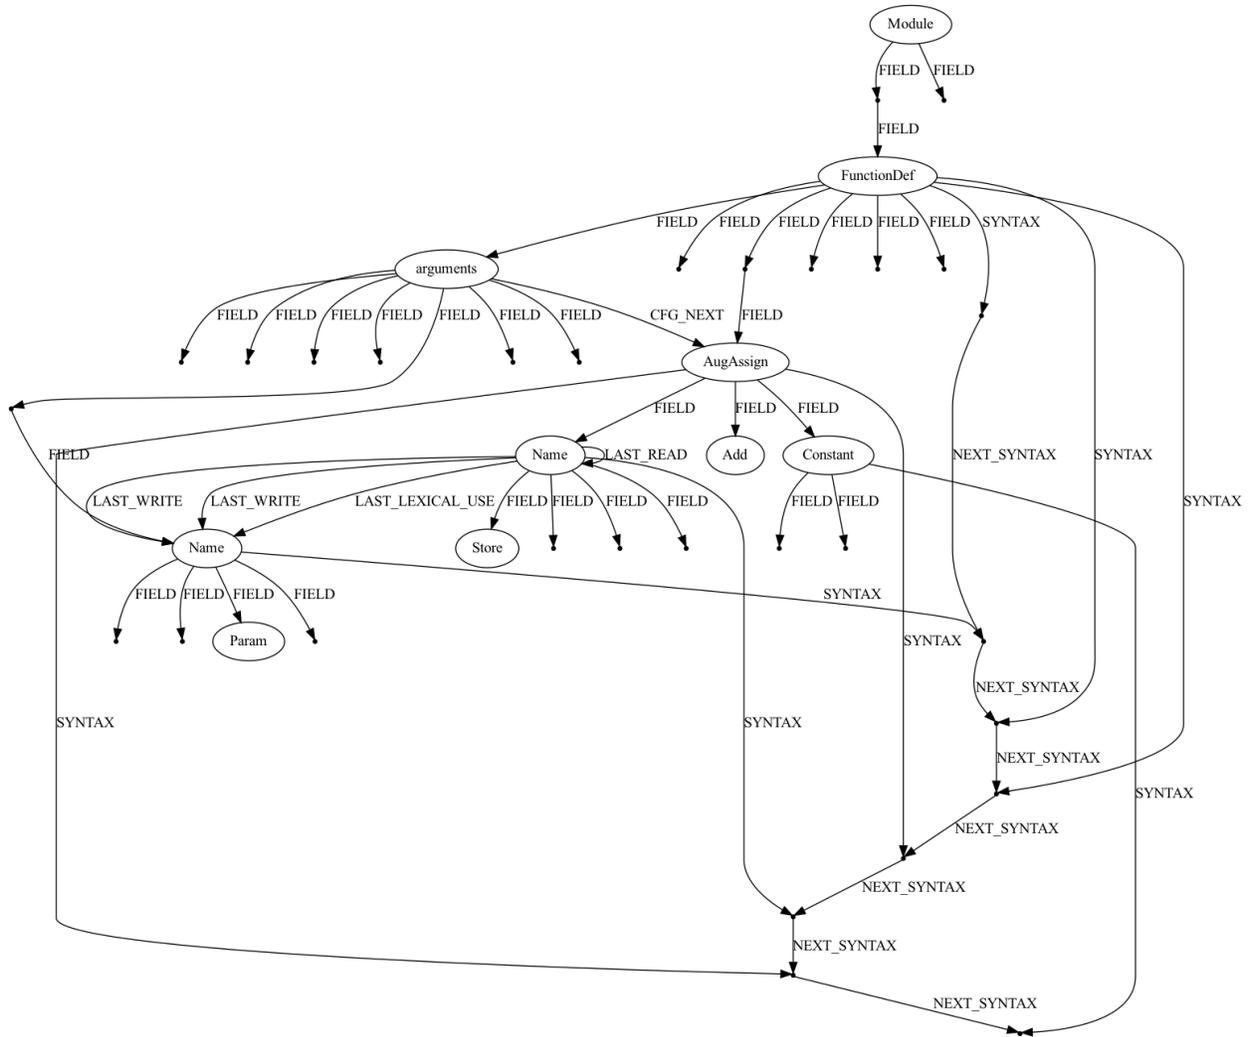


Figure 13: Program graph for Program #8 from Table 2.

## B Histograms of Program Graph Metrics

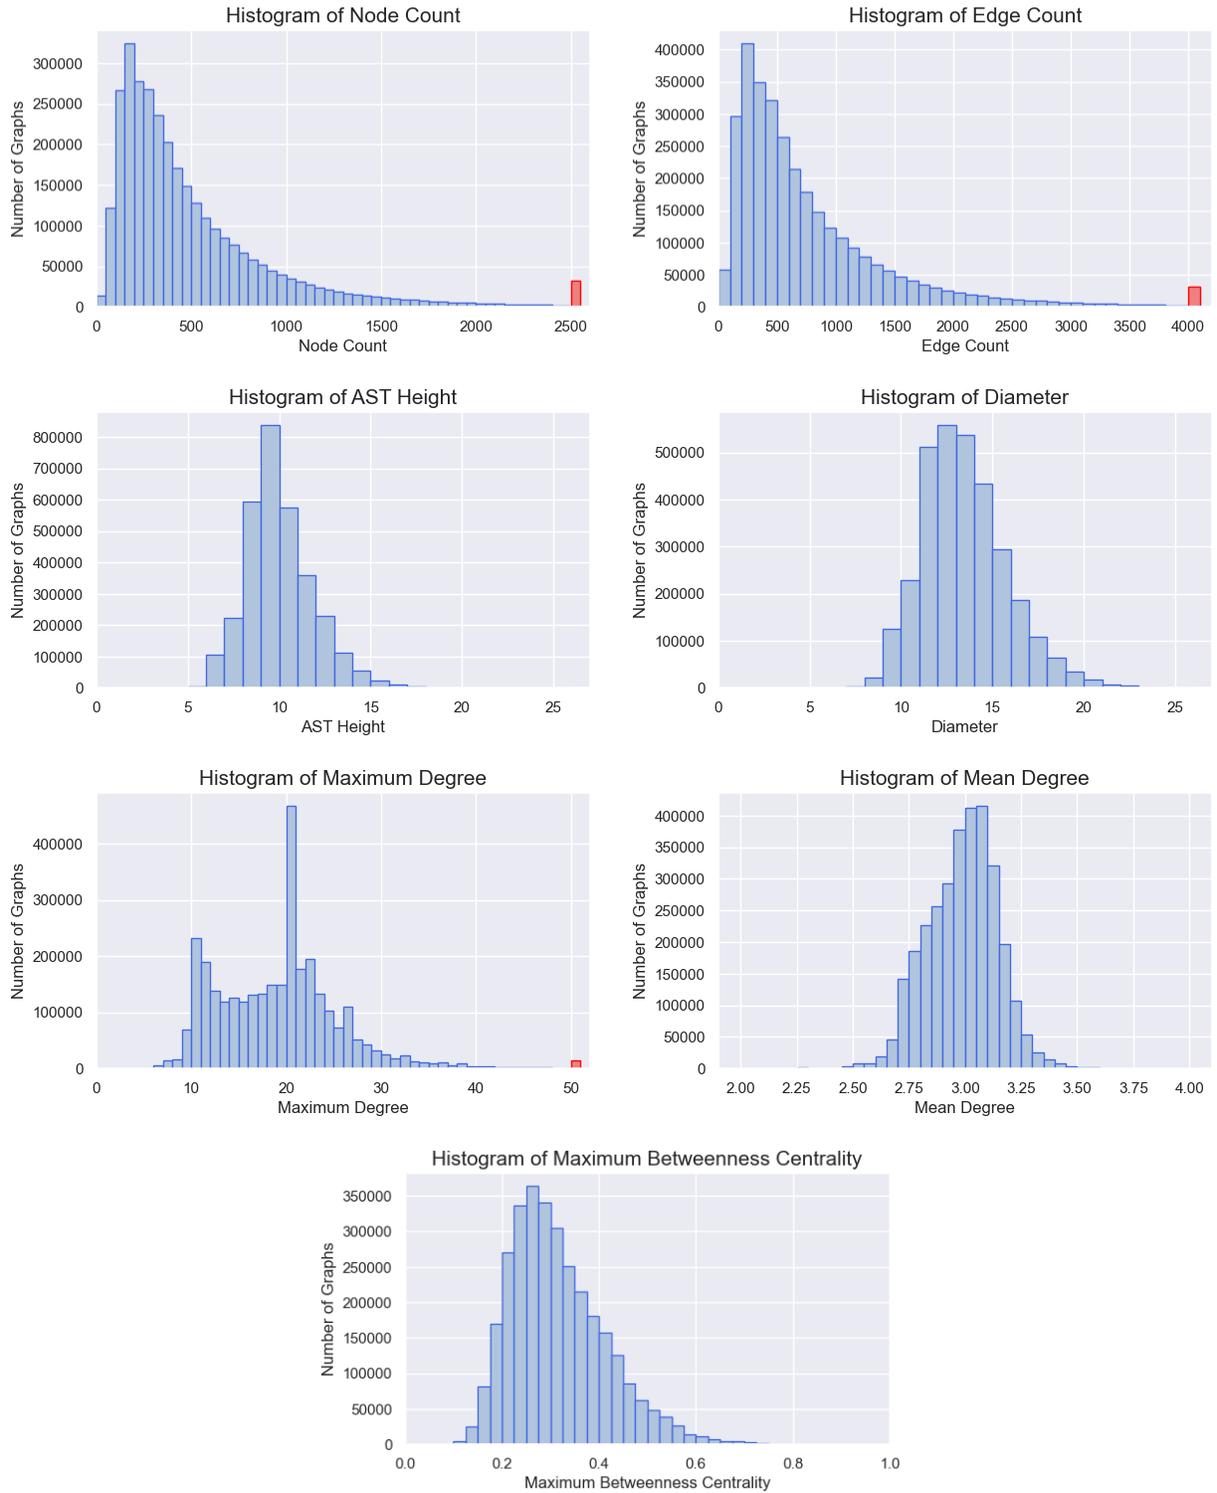


Figure 14: Histograms for various metrics of program graphs across the Project CodeNet dataset. Red bars include the program graphs for which the metric falls outside the range covered by the other bars.