

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

7-2021

On the generalizability of Neural Program Models with respect to semantic-preserving program transformations

Md Rafiqul Islam RABIN
University of Houston

Nghi D. Q. BUI
Singapore Management University, dqnbui@smu.edu.sg

Ke WANG
Visa Research

Yijun YU
Open University

Lingxiao JIANG
Singapore Management University, lxjiang@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Software Engineering Commons](#)

Citation

RABIN, Md Rafiqul Islam; BUI, Nghi D. Q.; WANG, Ke; YU, Yijun; and JIANG, Lingxiao. On the generalizability of Neural Program Models with respect to semantic-preserving program transformations. (2021).

Information and Software Technology. 135, 1-13.

Available at: https://ink.library.smu.edu.sg/sis_research/6044

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

On the Generalizability of Neural Program Models with respect to Semantic-Preserving Program Transformations

Md Rafiqul Islam Rabin^a, Nghi D. Q. Bui^c, Ke Wang^b, Yijun Yu^d, Lingxiao Jiang^c, Mohammad Amin Alipour^a

^a*University of Houston, United States.*

^b*Visa Research, United States.*

^c*Singapore Management University, Singapore.*

^d*The Open University, UK.*

Abstract

Context: With the prevalence of publicly available source code repositories to train deep neural network models, neural program models can do well in source code analysis tasks such as predicting method names in given programs that cannot be easily done by traditional program analysis techniques. Although such neural program models have been tested on various existing datasets, the extent to which they generalize to unforeseen source code is largely unknown. **Objective:** Since it is very challenging to test neural program models on all unforeseen programs, in this paper, we propose to evaluate the generalizability of neural program models with respect to semantic-preserving transformations: a generalizable neural program model should perform equally well on programs that are of the same semantics but of different lexical appearances and syntactical structures. **Method:** We compare the results of various neural program models for the method name prediction task on programs before and after automated semantic-preserving transformations. We use three Java datasets of different sizes and three state-of-the-art neural network models for code, namely *code2vec*, *code2seq*, and *GGNN*, to build nine such neural program models for evaluation. **Results:** Our results show that even with small semantically preserving changes to the programs, these neural program models often fail to generalize their performance. Our results also suggest that neural program models based on data and control dependencies in programs generalize better than neural program models based only on abstract syntax trees (ASTs). On the positive side, we observe that as the size of the training dataset grows and diversifies the generalizability of correct predictions produced by the neural program models can be improved too. **Conclusion:** Our results on the generalizability of neural program models provide insights to measure their limitations and

provide a stepping stone for their improvement.

Keywords: neural models, code representation, model evaluation, program transformation, generalizability

1. Introduction

Abundance of publicly available source code repositories has enabled a surge in data-driven approaches to programs analysis tasks. Those approaches aim to discover common programming patterns for various downstream applications [3] that are not easily achievable via traditional program analysis techniques, *e.g.*, prediction of data types in dynamically typed languages [23], detection of the variable naming issues [4], or repair of software defects [17]. The advent of deep neural networks has accelerated the innovation in this area and has greatly enhanced the performance of these approaches. The performance of deep neural networks in cognitive tasks such as method name prediction or variable naming has reached or exceeded the performance of other data-driven approaches. The performance of neural networks has encouraged researchers to increasingly adopt neural networks in program analysis tasks, giving rise to increasing uses of neural program models.

While the performance of neural program models continues to improve, the extent to which they can generalize to new, unseen programs is still unknown, even if the programs are in the same programming language. This problem is of more importance if we want to use them in downstream safety-critical tasks, such as malware detection and automated defect repair. This problem is particularly hard, as the interpretation of neural models that constitute the core reasoning engine of neural program models remains challenging—especially for the complex neural networks (*e.g.*, RNN) that are commonly used in the proposed neural program models.

A comprehensive understanding of the extent of generalizability of neural program models would help developers to know when to use data-driven approaches and when to resort to traditional deductive methods of program analysis. It would also help researchers to focus their efforts on devising new techniques to alleviate the shortcomings of existing neural program models. Lack of knowledge about the limits of neural program models may exaggerate their capability and cause careless applications of the neural program models on the domains that they are not suited for; or, spending time and efforts on developing neural program models while a traditional, more understandable technique can perform equally well or better.

Recently, we have seen a growing interest in the rigorous evaluation of neural program models. Wang and Christodorescu [50] compared the robustness of different program representations under compiler optimization transformations. They found that the program representations based on static code features are more sensitive to such changes than dynamic code features. Allamanis [1] evaluated the impact of code duplication in various neural program models and found that code duplication in the training and test datasets inflated the performance of almost all current neural program models. More recently, preliminary studies in this field started to emerge; *e.g.*, Rabin et al. [40] proposed the idea of testing neural program models using semantic-preserving transformations; Bui et al. [13] measured the impact of a specific code fragment by deleting it from the original source code; Zhang et al. [58] proposed a sampling approach to generate adversarial examples for code classification models; and Compton et al. [15] showed that the obfuscation of variable names makes a model on source code more robust with less bias towards variable names. Further, Yefet et al. [56] followed and proposed adversarial example generation for neural program models using prediction attribution [45]; Ramakrishnan et al. [41] increased robustness of neural representations of code by adding semantically equivalent programs to the training data; and Bielik and Vechev [11] proposed an approach for increasing the robustness of neural program models for type prediction based on finding prediction attribution, adversarial training, and refining source code representations. Although these studies share the similar ultimate goal of evaluating and improving the performance of neural program models with respect to unseen programs, there is still a lack of systematic quantifiable metrics to measure the extent to which the neural program models can generalize to unseen programs, and it would not be fair either to evaluate a neural program model against all possible unseen programs that it was not designed for.

Goal. In this paper, we attempt to understand the limits of *generalizability* of neural program models by comparing their behavior before and after semantic-preserving program transformations. That is, how the results of a neural program model generalize to a semantically-equivalent program. By limiting unseen programs to semantically equivalent ones and controlling the semantic-preserving program transformations, we are able to provide a fair, systematic, quantifiable metric for evaluating the generalizability of a neural program model.

In this paper, we report the results of a study on the generalizability of three highly-cited neural program models: *code2vec* [7], *code2seq* [6], and *GGNN* [19]. To evaluate their generalizability, we transform programs in the original datasets for testing to generate semantically-equivalent counter-

parts. We employ six semantic-preserving transformations that impact the structure of programs (*i.e.* abstract syntax trees) with varying degrees, ranging from common refactoring, *e.g.*, variable renaming, to more intrusive changes such as changing for-loops to while-loops.

Our results suggest that all neural program models evaluated in this study are sensitive to the semantic-preserving transformations; that is, the output of the neural program model would be different on transformed programs compared to its output on the original programs. This sensitivity remains an issue even in the cases of small changes to the programs, such as renaming variables or reordering independent statements in a basic block. Moreover, our results suggest that neural program models (*e.g.*, *GGNN*) that encode data and control dependencies in programs generalize better than the neural program models that are solely based on abstract syntax trees, and in most cases the generalizability of a neural program model can be improved with the growth in the size of training datasets.

The results of this study reveal that the generalizability of neural program models is still far from ideal and require more attention from the research community to devise more generalizable models of source code, or designing pre-processing techniques, *e.g.* canonicalizing program representations, to increase immunity of neural program models to such program transformations.

Compared to closely related work by Yefet et al. [56] and Ramakrishnan et al. [41] where their goals are adversarial code generation and increasing robustness of neural program models, this paper provides a complementary view to the evaluation of neural program models by focusing on the evaluation of generalizability of neural program models with a large number of transformations, and in-depth analysis of changes in their behavior on transformed programs. This paper also evaluates the impact of the size of datasets and programs on the generalizability of neural program models.

Contributions. This paper makes the following contributions.

- We introduce the notion of generalizability with respect to semantic-preserving transformations for neural program models.
- We perform a large-scale study to evaluate the generalizability of state-of-the-art neural program models. We also provide insights into the generalizability of existing neural program models and discuss their practical implications.
- We provide an in-depth analysis of changes in the prediction and evaluate the impact of the size of datasets and programs on the generalizability of neural program models.

<pre> public int compareTo(ApplicationAttemptId other) { int compareAppIds = this.getApplicationId() .compareTo(other.getApplicationId()); if (compareAppIds == 0) { return this.getAttemptId() - other. getAttemptId(); } else { return compareAppIds; } } </pre>	<pre> public int compareTo(ApplicationAttemptId var0) { int compareAppIds = this.getApplicationId() .compareTo(var0.getApplicationId()); if (compareAppIds == 0) { return this.getAttemptId() - var0. getAttemptId(); } else { return compareAppIds; } } </pre>
Prediction before transformation: compareTo	Prediction after transformation: getCount

Figure 1: Variable Renaming on java-small/test/hadoop/ApplicationAttemptId.java file.

2. Motivating Example & Definition

We use *code2vec* [7] for exposition in this section. The *code2vec* [7] is a recent, highly-cited (200+ citations as of Nov. 2020) neural program model that predicts the name of a Java method given the body of the method. Such a neural program model can assist developers in classification of methods, code similarity detection, and code search.

Figure 1 shows two semantically-identical methods that implement `compareTo` functionality. The only difference between them is in the name of one of the variables. The left snippet in Figure 1 uses `other`, while the code on the right uses `var0`¹. However, the *code2vec* outputs, *i.e.*, predictions, on these semantically equivalent programs are drastically different. *code2vec* predicts the snippet on the left to be `compareTo` function, and the function on the right to be `getCount`. It seems that the predictions of *code2vec* rely much on the identifier names (*e.g.*, `other`). This reliance would make *code2vec* susceptible to a common refactoring such as variable renaming, and would make it not generalize to the code snippets that are semantically the same, but are different syntactically, even under common transformations.

Lack of generalizability would lead to distrust in the neural program models and hamper their wider adoption and application. If such neural program models were to be deployed in the problem settings wherein higher levels of generalizability are required, *e.g.*, malware detection and bug repair, it would be much better for the neural program models to demonstrate a high level of generalizability with respect to certain metrics.

Generalizability. We define generalizability as the capability of a neural program model to return

¹`var0` is not an uncommon identifier name in Java as it appears in the training vocabulary of the datasets. At the time of writing, a search on the GitHub returns more than 75K Java classes that use this identifier.

the same results under semantic-preserving transformations.

In this paper, we differentiate *generalizability* from the term *robustness* that is commonly used in the neural network literature [46] for two main reasons. First, robustness is usually defined in the face of adversarial examples that have security implications, while we do not generate adversarial examples. Second, robustness implies *imperceptible* differences in the two focal inputs (*e.g.*, minor pixel changes in two images) that are hard to attain in a sparse domain such as source code; the program transformations used in this paper often lead to perceptible changes of textual appearances and syntactic structures in program code. We also note that our definition of generalizability differs from what is used in [28] that evaluates the usefulness of a neural program model in various downstream tasks.

Together with clearly defined semantic-preserving program transformations and their change impact on the prediction results of neural program models (cf. Section 4), we aim to provide a systematic quantifiable way to measure the generalizability of neural program models, and thus shed lights on their capabilities and limits for future improvements. With the extensibility of the program transformations and the measurements of their change impact, our evaluation approach may also be extended to measure the generalizability of neural program models more comprehensively in the near future.

3. Background

Most neural program models use neural network classifiers in their core components that take a code snippet or a whole program as an input, and make predictions about some of its characteristics; *e.g.*, a bug prediction classifier that predicts the buggy-ness of statements in the input program.

Performance of a neural program model depends on three main factors: quality of data (*i.e.*, source code for this study), the representation of data for the neural network, and the neural network characteristics and its training parameters.

Quality of the data is concerned with the representativeness of data, and proper cleaning and preprocessing of the data. Currently, most studies use open-source projects usually in mainstream programming languages, *e.g.*, C#, Java, C, or JavaScript. The available datasets for these tasks are still very immature and not standardized, and their quality is somewhat unknown. For example, a recent study by Allamanis [1] showed that virtually all available datasets suffer from code duplication that can greatly impact the performance of neural program models.

The second factor affecting the performance of neural program models is source code representations. Since neural networks need to take vectors of numbers as direct inputs, source code embeddings are used to produce a vector representation of source code. The representation determines which program features to include and how they should be represented in the vector embeddings. The representations can be broadly categorized into two categories: static and dynamic. Static program representations consider only the features that can be extracted from parsing texts of the programs, while dynamic representations include some features pertaining to the real executions of the programs.

The third factor impacting the performance of a neural program model is the characteristics—*e.g.*, type, topology, and hyper-parameters—of the neural networks it uses. There are numerous choices of network architectures each with different properties. Currently, the class of recurrent neural networks (*e.g.*, LSTM) and graph neural networks are among the most popular architectures in neural program models [7, 6, 19].

4. Evaluation Approach

Our approach for evaluating neural program models relies on a metamorphic relation that states: *the outputs of a neural program model should not differ on semantically-equivalent programs*. To this end, the evaluation approach is divided into two main steps: (1) generating new programs using semantic-preserving transformations, and (2) comparing the outputs of a neural program model before and after the transformations to compute generalizability metrics. We describe these steps in the rest of this section.

4.1. Target Downstream Task

We use the method name prediction task [2, 5] in this work to evaluate the generalizability of neural program models. The goal of the task is to predict the name of a method given the body of the method. This task has several applications such as code search [35], code summarization [5], and code analogies [7]. Figure 1 depicts an example of this task wherein neural program models are given a method body and return candidate names for the method body, *i.e.*, `compareTo` and `getCount`. This task has been used as the downstream task to evaluate several state-of-the-art neural program models [4, 7, 6].

4.2. Transformations

In this work, we only evaluate neural program models that take a method body as their input, therefore, we use the following set of transformations that are applicable to method-level code to generate semantically-equivalent methods. This set includes transformations ranging from common refactoring like variable renaming to more intrusive ones like loop exchange. The goal is to evaluate the generalizability of neural program models under a wide range of semantic-preserving changes to the structure of a method.

- **Variable Renaming (VN)** is a refactoring that renames the name of a variable in a method. The new name of the variable will be in the form of `varN` for a value of `N` such that `N` has not been defined in the scope. VN is a widely-used refactoring for methods.
- **Permute Statement (PS)** swaps two independent statements (i.e., with no data or control dependence) in a basic block of a method.
- **Unused Statement (UN)** inserts an unused string declaration to a randomly selected basic block in a method. Unused variables in methods are a common malpractice by developers.
- **Loop Exchange (LX)** replaces `for` loops with `while` loops or vice versa.
- **Switch to If (SF)** replaces a `switch` statement in a method with an equivalent `if` statement.
- **Boolean Exchange (BX)** switches the value of a boolean variable in a method from `true` to `false` or vice versa, and propagates this change in the method to ensure a semantic equivalence of the transformed method with the original method.

Note that each transformation has different impact on the structure of methods as follows.

- The *Variable Renaming* transformation only changes the terminal values and does not affect the structure of an AST.
- The *Permute Statement* transformation does not change the nodes, rather it only reorders two subtrees in an AST.
- The *Unused Statement* transformation adds a few nodes into an AST, which increases the number of paths in the AST.

- The *Loop Exchange* transformation extensively impacts an AST by removing and inserting nodes.
- The *Switch to If* transformation also impacts the AST of a method substantially by removing and inserting nodes.
- The *Boolean Exchange* transformation alters the value of `true` or `false` and modifies the structure of an AST by removing or inserting unary-not nodes.

4.3. Generalizability Metrics

In this study, we define a few metrics to measure different results of a neural program model for transformed programs and thus to quantify the generalizability of the neural program model.

Specifically, suppose M denotes a set of methods, given a semantic-preserving program transformation T that takes a method and creates a set $M' = \bigcup_{m \in M} T(m)$ of transformed methods, and a neural program model $NPM : M \rightarrow L$, where L denotes a set of labels, maps methods to labels. We evaluate the generalizability of NPM with respect to the transformation T , by comparing $NPM(m)$ and $NPM(m')$ for $m' \in T(m)$ for $m \in M$. Ideally, the neural program model should produce the same results on both m and m' , that is $NPM(m) = NPM(m')$. We define the following metrics.

Prediction Change Percentage. We compute the prediction change percentage as follows:

$$PCP = \frac{|\{m' \in M' | NPM(m) \neq NPM(m')\}|}{|\{m' \in M'\}|} * 100. \quad (1)$$

The lower values of PCP for NPM would suggest higher a degree of its generalizability with respect to the transformation.

Types of Changes. Considering that the correctness of predicted labels of the NPM , five types of changes can happen:

- (1) a correct prediction remains correct after the transformation,
- (2) a correct prediction changes to a wrong prediction after the transformation,
- (3) a wrong predicted label remains the same wrong label after the transformation,
- (4) a wrong prediction changes to a correct prediction after the transformation,
- (5) a wrong predicted label changes into a different, yet still wrong label after the transformation.

We use the following five metrics to denote the proportion of each of these cases in the experiments. **CCP**, **CWP**, **WWSP**, **WCP**, and **WWDP** respectively denote the percentage of correct predictions that stay correct, the percentage of correct predictions that become wrong, the percentage of wrong predictions that stay to the same wrong prediction after the transformation, the percentage of wrong predictions that become correct, and the percentage of wrong predictions that change to a different wrong prediction after the transformation.

Precision, Recall, and F₁-Score. We also use the traditional sub-token metrics (precision, recall and F₁-score) as commonly used in the literature for the method name prediction task [7, 6] in this generalizability study. Suppose, *tp* denotes the number of true positive sub-tokens, *fp* denotes the number of false positive sub-tokens, and *fn* denotes the number of false negative sub-tokens in the predicted method names.

- *Precision* indicates the percentage of predicted sub-tokens that are true positives. It is the ratio of the correctly predicted positive sub-tokens to the total number of predicted positive sub-tokens: $Precision = \frac{tp}{tp + fp}$
- *Recall* indicates the percentage of true positive sub-tokens that are correctly predicted. It is the ratio of the correctly predicted positive sub-tokens to the total number of sub-tokens in actual method names: $Recall = \frac{tp}{tp + fn}$
- *F1-Score* is the harmonic mean of precision (P) and recall (R):

$$F_1\text{-Score} = \frac{2}{P^{-1} + R^{-1}} = 2 \cdot \frac{P \cdot R}{P + R}$$

For example, a predicted name `result_compute` has two sub-tokens `result` and `compute`, and is considered as an exact match of the ground-truth name `computeResult` which also has the same two sub-tokens (ignoring the case and the ordering of the tokens). Similarly, a predicted name `compute` has 100% precision but only 50% recall with respect to the same ground truth, and `compute_model_result` has 100% recall but only 67% precision.

5. Experimental Setting

5.1. Subject Neural Program Models

The task of method name prediction [5] has attracted some attention recently. We use three neural program models that use different code representations and neural network characteristics for the task: *code2vec* [7], *code2seq* [6], and *GGNN* [19].

code2vec [7] uses a bag of AST paths to model source code. Each path consists of a pair of terminal nodes and the corresponding path between them in the AST. Each path, along with source and destination terminals, is mapped into its vector embeddings which are learned jointly with other network parameters during training. The separate vectors of each path-context are then concatenated to a single context vector using a fully connected layer which is learned during training with the network. An attention vector is also learned with the network; it is used to score each path-context and aggregate multiple path-contexts to a single code vector representing a method body. After that, the model predicts the probability of each target method name given the code vector of the method body via a softmax-normalization between the code vector and each of the embeddings of all possible target method names.

While *code2vec* uses monolithic path embeddings and only generates a single label at a time, the *code2seq* [6] model uses an encoder-decoder architecture to encode paths node-by-node and generate labels as sequences at each step. In *code2seq*, the encoder represents a method body as a set of AST paths where each path is compressed to a fixed-length vector using a bi-directional LSTM which encodes paths node-by-node. The decoder uses attentions to select relevant paths while decoding, and predicts sub-tokens of a target sequence at each step when generating the method name.

In *GGNN* [19], a variety of semantic edges are added into the AST of a method body to construct a graph, and the Gated Graph Neural Network (GGNN) is applied to encode such graphs [4]. The initial embedding for a node of the graph is the concatenation between the node type embedding and node token embedding. Then a fixed number of message passing steps are applied for a node to aggregate the embeddings of its neighbors. The output of the GGNN encoder is then fed into a bi-directional LSTM decoder to generate the method name as a language model of sub-tokens [19].

5.2. Datasets

We have used the *code2seq* dataset for training neural program models for the study. There are three Java datasets based on the GitHub projects: JAVA-SMALL, JAVA-MED, and JAVA-LARGE.

- **JAVA-SMALL:** This dataset contains 9 Java projects for training, 1 for validation and 1 for testing. Overall, it contains about 700K methods. The compressed size is about 366MB and the extracted size is about 1.9GB.
- **JAVA-MED:** This dataset contains 800 Java projects for training, 100 for validation and 100 for testing. Overall, it contains about 4M methods. The compressed size is about 1.8GB and

Table 1: Performance of trained models for method name prediction in the testing dataset.

Model	Dataset	# Original methods in the testing dataset	Precision	Recall	F_1 -Score
<i>code2vec</i>	JAVA-SMALL	44426	28.36	22.37	25.01
	JAVA-MED	351628	42.55	30.85	35.76
	JAVA-LARGE	370930	45.17	32.28	37.65
<i>code2seq</i>	JAVA-SMALL	44426	46.30	38.81	42.23
	JAVA-MED	351628	59.94	48.03	53.33
	JAVA-LARGE	370930	64.03	55.02	59.19
<i>GGNN</i>	JAVA-SMALL	44426	49.12	47.18	48.59
	JAVA-MED	351628	58.30	47.49	52.34
	JAVA-LARGE	370930	60.76	50.32	55.23

the extracted size is about 9.3GB.

- JAVA-LARGE: This dataset contains 9000 Java projects for training, 200 for validation and 300 for testing. Overall, it contains about 16M methods. The compressed size is about 7.2GB and the extracted size is about 37GB.

5.3. Training Models per Datasets

The authors of *code2vec* and *code2seq* have made the source code public for training and evaluating their models. For *GGNN*, the implementation of the network is available but the code graph generation is not; so we re-implement the step to generate graphs. We use the parser SrcSlice², an extension of SrcML³, to produce data dependency edges among AST nodes for training *GGNN*.

We train each model for the method name prediction task with the configurations described in their original papers on each of the three aforementioned datasets, and thus construct three *code2vec*, three *code2seq*, and three *GGNN* neural program models. Similar to the state-of-the-art approaches, i.e. [7, 6], we train models on the training set, tune on the validation set for maximizing F_1 -score, and finally report results on the unseen testing set. Table 1 summarizes the performance of trained models for method name prediction on the testing set. While the performance of our trained models for *code2seq* is on par with to the ones reported in the corresponding paper [6], the performance of *code2vec* did not reach the performance reported in [7], due to the differences in the dataset. However, the performance of our trained *code2vec* models is similar to the one reported in [6]. For *GGNN*, the performance is reasonably different from what were reported in [19] for

²<https://github.com/srcML/srcSlice>

³<https://www.srcml.org/>, +400 node types for supporting multiple programming languages

mainly three reasons: (1) the ASTs produced by our parser are different, (2) the extraction of some types of semantic edges proposed in [19] requires expensive analysis of the methods; therefore, we implemented and included only a subset (seven out of ten) of semantic edges into the ASTs when constructing the graphs, and (3) the datasets are different.

5.4. Population of Transformed Programs

We have used our own tool based on the `JavaParser`⁴ library to transform Java methods. Henceforth we use terms program and method interchangeably. Two authors were involved in the implementation, testing and code review. We have performed manual inspection of sample transformed programs to ensure correctness of the transformations.

We have applied the applicable transformations to the methods available in the *testing* data of the three datasets mentioned in Section 5.2. The number of original methods in our study is 1,415,116,⁵. Overall, the number of original methods with incorrect predictions is, on average, 2.8 times higher than the number of methods with correct predictions.

We create a set of *single-place* transformed programs (Table 2) by applying transformations to each eligible location in methods separately resulting in 2,822,810 transformed methods; e.g., if a method has three eligible locations for a transformation, we would generate three distinct methods by transforming each individual location separately. The types and number of applicable transformations vary from a method to another. Therefore, in our approach, different methods, based on the language features that they use, produce a different number of transformed programs. In total, the number of transformed programs generated from the programs with incorrect initial predictions is much higher (4.2x and higher) than the number of transformed programs generated from the programs with correct initial predictions, which may suggest that programs with correct predictions may be smaller and simpler.

Artifacts. The source code of the transformation tool and transformed datasets used in this paper are publicly available at <https://github.com/mdrafiqulrabin/tnpa-generalizability/>.

5.5. Research Questions

In this paper, we seek to answer the following research questions.

⁴<https://github.com/javaparser/javaparser>

⁵This total number is different from the numbers in Table 1 because a method in the testing dataset may contain code elements eligible for multiple types of transformations and be counted multiple times.

Table 2: Prediction Change Percentage (PCP) across all models, datasets, and transformations.

Transformation	Dataset	# Original methods	# Transformed methods	Prediction change (%) (PCP)		
				<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>
<i>Variable Renaming</i>	JAVA-SMALL	31113	123123	54.92	57.16	28.17
	JAVA-MED	235961	771208	46.55	48.75	35.96
	JAVA-LARGE	252725	916565	42.06	47.04	31.92
	Weighted Average =			44.85	48.46	33.39
<i>Boolean Exchange</i>	JAVA-SMALL	1158	1519	53.85	54.31	29.37
	JAVA-MED	6407	8840	50.35	44.71	33.74
	JAVA-LARGE	8868	12107	47.80	51.43	31.98
	Weighted Average =			49.21	48.98	32.50
<i>Loop Exchange</i>	JAVA-SMALL	3699	5160	59.38	52.54	31.66
	JAVA-MED	17107	23533	62.77	45.29	36.67
	JAVA-LARGE	35565	49665	46.52	42.51	31.75
	Weighted Average =			52.25	44.01	33.22
<i>Switch to If</i>	JAVA-SMALL	246	259	68.73	61.78	31.45
	JAVA-MED	3312	3839	59.91	41.60	43.73
	JAVA-LARGE	10478	11165	30.33	29.08	45.50
	Weighted Average =			38.42	32.78	44.82
<i>Permute Statement</i>	JAVA-SMALL	3397	9169	72.80	57.32	26.36
	JAVA-MED	16150	44711	65.44	42.64	34.09
	JAVA-LARGE	21956	74973	64.38	41.93	26.32
	Weighted Average =			65.35	43.27	29.02
<i>Unused Statement</i>	JAVA-SMALL	44426	44426	39.97	45.60	28.34
	JAVA-MED	351621	351621	35.80	40.25	42.79
	JAVA-LARGE	370927	370927	31.21	37.44	35.67
	Weighted Average =			33.82	39.20	38.51

RQ1 How do the transformations impact the predictions of neural program models in the single-place transformed dataset?

RQ2 When do the transformations affect neural program models the most?

RQ3 How does the method length impact the generalizability of neural program models?

RQ4 What are the trends in types of changes?

RQ5 How do the transformations affect the precision, recall and F_1 -score of the neural program models?

6. Results

6.1. RQ1: Impact of Transformations on the Predictions of Neural Program Models

Table 2 shows the prediction change percentage (PCP) of the neural program models for each transformation and dataset. In this table, “# Original methods” denotes the number of methods eligible for the corresponding transformation, “# Transformed methods” denotes the number of

methods generated as the result of applying the corresponding transformations on the original methods, and “Prediction change (%)” denotes PCP as defined in Section 4. “Weighted Average” provides the weighted average of PCP for each transformation and neural models. The bold values in the Table 2 highlight the minimum value of PCP for the transformations. Since a transformation can be applied in more than one place separately in a method body, the number of transformed methods can be larger than the number of original methods.

As Table 2 depicts, all neural program models are likely to susceptible to semantic-equivalent transformations; however, the impact of transformations on PCP differs among different neural networks and datasets. Overall, *GGNN* seems less prone to prediction changes; in 14 out of 18 cases, PCP in *GGNN* is significantly less than *code2vec* and *code2seq*. Moreover, in four out of six transformations, the weighted average of PCP for *GGNN* is lower than the rest.

Observation 1: In most cases, *GGNN* seems less susceptible to prediction changes under semantic-preserving transformations, compared to *code2vec* and *code2seq*.

Within *code2vec*, *code2seq*, and *GGNN*, the PCP trend varies for different transformations and datasets. *code2vec* is comparatively most sensitive to *Permute Statement* on all datasets. On the other hand, *code2seq* is most vulnerable to *Switch to If* in JAVA-SMALL, *Variable Renaming* in JAVA-MED, and *Boolean Exchange* in JAVA-LARGE. In *GGNN*, *Switch to If* is the most powerful transformation on all datasets. In most cases, for *code2vec* and *code2seq*, the PCP for *Unused Statement* is comparatively less than the other transformations, except for *code2seq* in JAVA-LARGE where *Switch to If* is less sensitive. In *GGNN*, *Permute Statement* is a comparatively less powerful transformation than others on all datasets. Overall, based on the weighted average, it is likely that, *code2vec* is most sensitive to *Permute Statement* and least sensitive to *Unused Statement*, *code2seq* is most sensitive to *Boolean Exchange* and least sensitive to *Switch to If*, and *GGNN* is most sensitive to *Switch to If* and least sensitive to *Permute Statement*.

Based on the weighted average, *GGNN* performs worst for *Switch to If* and *Unused Statement* transformations. These two transformations add some additional nodes and paths in the AST. For *code2vec* and *code2seq*, if models give less attention to those new paths, then the change can be less effective. However, *GGNN* works by using a message passing mechanism among the nodes with a limited number of passing steps. In *Unused Statement*, because there is some irrelevant information added into the code, the passing steps in *GGNN* can capture this information and ignore other useful information, thus having a strong impact on the prediction results. In *Switch to If*, because

the structure of the AST is modified by adding and removing nodes, and *GGNN* is a node-based method, i.e., combining node information with message passing, thus the *GGNN* can sensitive to node modification in the AST for *Switch to If*.

Table 2 also supports that, in most cases, *Permute Statement* is more powerful than *Variable Renaming* in *code2vec* model whereas *Variable Renaming* is more powerful than *Permute Statement* in *code2seq* model. This is probably caused by the real-value embeddings of AST paths are different for *code2vec* and *code2seq*. In *code2vec*, an embedding matrix is initialized randomly for paths and learned during training, that contains rows that are mapped to each of the AST paths. On the other hand, in *code2seq*, each node of a path comes from a learned embedding matrix, and then a bi-directional LSTM is used to encode each of the AST paths separately. The bi-directional LSTM reads the path once from beginning to the end (as original order) and once from end to beginning (in reverse order). Therefore, the order changes by *Permute Statement* may become less sensitive to *code2seq* than *code2vec*.

Another observation is that, in most cases of *code2vec* and *code2seq*, the PCP of the transformations in JAVA-SMALL is high, and it is significantly lower on larger datasets, i.e., JAVA-MED, and JAVA-LARGE. In *GGNN*, the PCP of the transformations shows a different trend: lowest in JAVA-SMALL, in most cases, and highest in JAVA-MED.

Observation 2: In most cases, the effect of prediction change for *code2vec* and *code2seq* is reduced as the dataset size increases, compared to *GGNN*.

6.2. RQ2: When Transformations Affect Neural Program Models the Most?

6.2.1. Single-place transformation vs. All-place transformation

In our analysis, thus far, if a program has multiple candidates for a transformation, say n candidates, for transformation, we only apply them one at the time and end up with n distinct transformed programs. We call this *single-place* transformation. Alternatively, we can apply the transformations to *all* candidate locations in the program simultaneously to create only one transformed program. We call this *all-place* transformation. We evaluate the generalizability of neural program models under all-place transformation for the following transformations: *Variable Renaming*, *Boolean Exchange*, *Loop Exchange*, and *Switch to If*. Note that the all-place transformation is not applicable to *Permute Statement* and *Unused Statement* transformations, as we apply the *Permute Statement* on a pair of statements and the *Unused Statement* on a random block.




Figure 2: Prediction Change Percentage (PCP) in single-place vs. all-place transformations.

Figure 2 compares the impact of single-place transformation and all-place transformation on the prediction changes in all neural program models. For the *code2vec* model, the percentage of prediction change for the all-place transformation is higher than the single-place transformation by a good margin for all the cases. Similarly, for the *code2seq* model, the percentage of prediction change for the all-place transformation is higher than the single-place transformation by a good margin except for the case (*Switch to If*, JAVA-SMALL). After a closer examination of JAVA-SMALL dataset and *Switch to If* transformation, we observe that the number of transformed methods for all-place is only 13, which is too low to provide comparative insight. For the *GGNN* model, the difference between all-place transformation and single-place transformation is relatively very small compared to the *code2vec* and *code2seq* models. Even for (*Boolean Exchange*, JAVA-SMALL + JAVA-MED), (*Loop Exchange*, JAVA-MED + JAVA-LARGE), and (*Switch to If*, JAVA-LARGE), the percentage of prediction changes for the single-place transformation is higher than the all-place transformation. The results may suggest that the performance of *GGNN* under single-place transformations and all-place transformations is almost consistent.

Observation 3: While all-place transformations are more likely to induce prediction changes in *code2vec* and *code2seq* than single-place transformations, the performance of *GGNN* remains relatively similar under both types of transformations.

6.2.2. Correctly predicted methods vs. Incorrectly predicted methods

We also evaluate the generalizability of neural program models under correctly and incorrectly predicted methods. Figure 3 compares the impact of correctly predicted methods and incorrectly predicted methods on the prediction changes in all neural program models. In the *code2vec* model, the percentage of changes in predictions after transformation in the correctly predicted methods

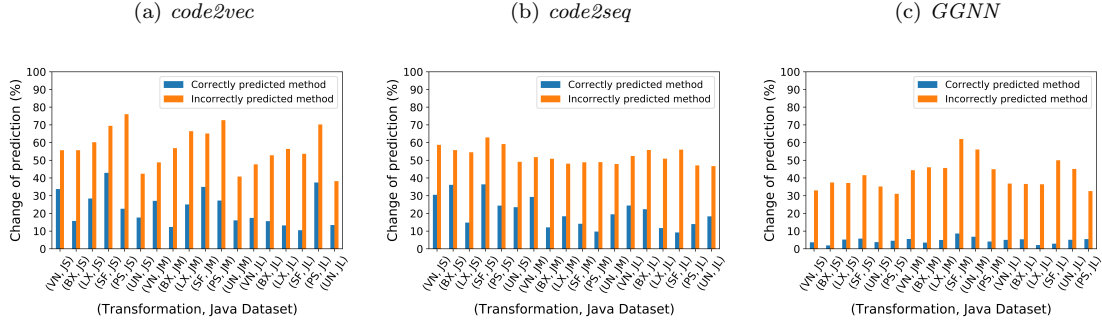


Figure 3: Prediction Change Percentage (PCP) on correctly vs. incorrectly predicted methods.

ranges from 10.45% to 42.86%, while, in the incorrectly predicted methods, a larger portion of transformations, 38.18% to 76.00%, change the prediction of *code2vec*. Similarly, in the *code2seq* model, the percentage of changes in predictions after transformation in the correctly predicted methods and in the incorrectly predicted methods ranges from 9.19% to 36.36% and 46.66% to 62.90%, respectively. However, in the *GGNN* model, while the percentage of changes in predictions after transformation on the correctly predicted methods ranges from 1.90% to 8.58%, the percentages range from 31.05% to 62.01% in the incorrectly predicted methods.

Observation 4: It is likely that *GGNN* is more stable than *code2vec* and *code2seq* in the originally correct methods, and the changes in prediction happen more frequently in the originally incorrect methods for all models.

6.2.3. The Effect of $X\%$ -Transformation

In this section, we evaluate the generalizability of neural program models under $X\%$ -transformation for the following transformations: *Variable Renaming*, *Boolean Exchange*, *Loop Exchange*, and *Switch to If*. Note that we apply the *Permute Statement* on a pair of independent statements and the *Unused Statement* on one randomly selected block in a method, and thus $X\%$ -transformation is not suitable for them. If a transformation t is applicable to n locations in a method body, $X\%$ -transformation randomly picks $\lfloor \frac{n \cdot X}{100} \rfloor$ of those locations and applies t to create a new transformed program. The number of all $X\%$ -transformed programs grows exponentially with the number of locations; therefore, to manage the complexity, in $X\%$ -transformation we randomly pick the locations in a method body, instead of considering all possible combinations, to create transformed programs. We study the $X\%$ -transformation with $X = \{25, 50, 75\}$. For each transformation t , we first create a dataset d_t^4 that contains methods with four or more possible locations, so that

Table 3: The PCP for $X\%$ -transformations across different datasets and models.

Dataset	Transformation	# Transformed methods	25% Transformation			50% Transformation			75% Transformation		
			<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>
JAVA-SMALL	<i>Variable Renaming</i>	15937	63.29	54.36	29.56	71.88	65.89	29.87	75.18	70.57	30.36
	<i>Boolean Exchange</i>	75	80.00	63.00	37.70	79.67	64.67	36.07	79.66	64.66	32.79
	<i>Loop Exchange</i>	302	81.95	65.90	32.44	81.87	65.98	32.44	81.38	65.56	34.73
	<i>Switch to If</i>	0	-	-	-	-	-	-	-	-	-
JAVA-MED	<i>Variable Renaming</i>	101003	54.07	46.51	37.18	62.65	57.77	38.91	66.51	63.20	37.88
	<i>Boolean Exchange</i>	428	69.66	48.20	31.37	70.50	48.60	30.97	70.91	47.49	31.17
	<i>Loop Exchange</i>	1292	86.01	55.11	28.72	86.11	57.37	26.29	85.37	57.62	29.01
	<i>Switch to If</i>	98	82.91	43.62	55.42	84.44	44.90	50.00	89.03	45.16	57.08
JAVA-LARGE	<i>Variable Renaming</i>	114748	45.62	43.23	33.37	53.32	53.99	35.75	56.98	58.83	34.59
	<i>Boolean Exchange</i>	642	71.81	59.03	28.18	71.09	62.76	31.76	71.93	62.03	30.40
	<i>Loop Exchange</i>	2899	79.77	56.00	25.96	79.02	56.79	27.41	78.57	56.92	27.05
	<i>Switch to If</i>	125	69.00	56.00	34.78	73.60	56.40	32.36	73.60	56.60	34.95

the transformation t is applicable to each method for each $X = \{25, 50, 75\}$, e.g., methods with at least four variables for *Variable Renaming*. To account for the randomness, we run each setting five times and report the average results, except for the *GGNN* that we run only three times due to the longer graph construction and processing time.

Table 3 shows the results of the $X\%$ -transformations. In each $X\%$ -transformation, *GGNN* has a much lower PCP value than the *code2vec* and *code2seq* models for all the transformations across the three Java datasets. Moreover, in *GGNN* models, the differences of PCP under different X are relatively small (mostly a few percentage points) and do not yield a clear trend. On the other hand, in *code2vec* and *code2seq* models, with the *Variable Renaming* transformation, the PCP tends to increase as X grows, but with other transformations except *Variable Renaming*, the PCP shows modest changes only. Note that in $X\%$ -transformation, compared to *Variable Renaming*, the numbers of transformed programs for other transformations are much lower, which might be too low to provide statistical significance or comparative insights.

Observation 5: The performance of *GGNN* in terms of PCP remains similar in all cases under $X\%$ -transformation, but the PCP of *code2vec* and *code2seq* for *Variable Renaming* increases as X grows.

6.3. RQ3: Impact of Method Length on Generalizability

An important metric of interest might be the generalizability in terms of the number of statements in the methods. Figure 4 depicts the relation between the length of methods and the prediction changes percentage (*i.e.*, PCP) in the single-place transformed data. In the figure, the

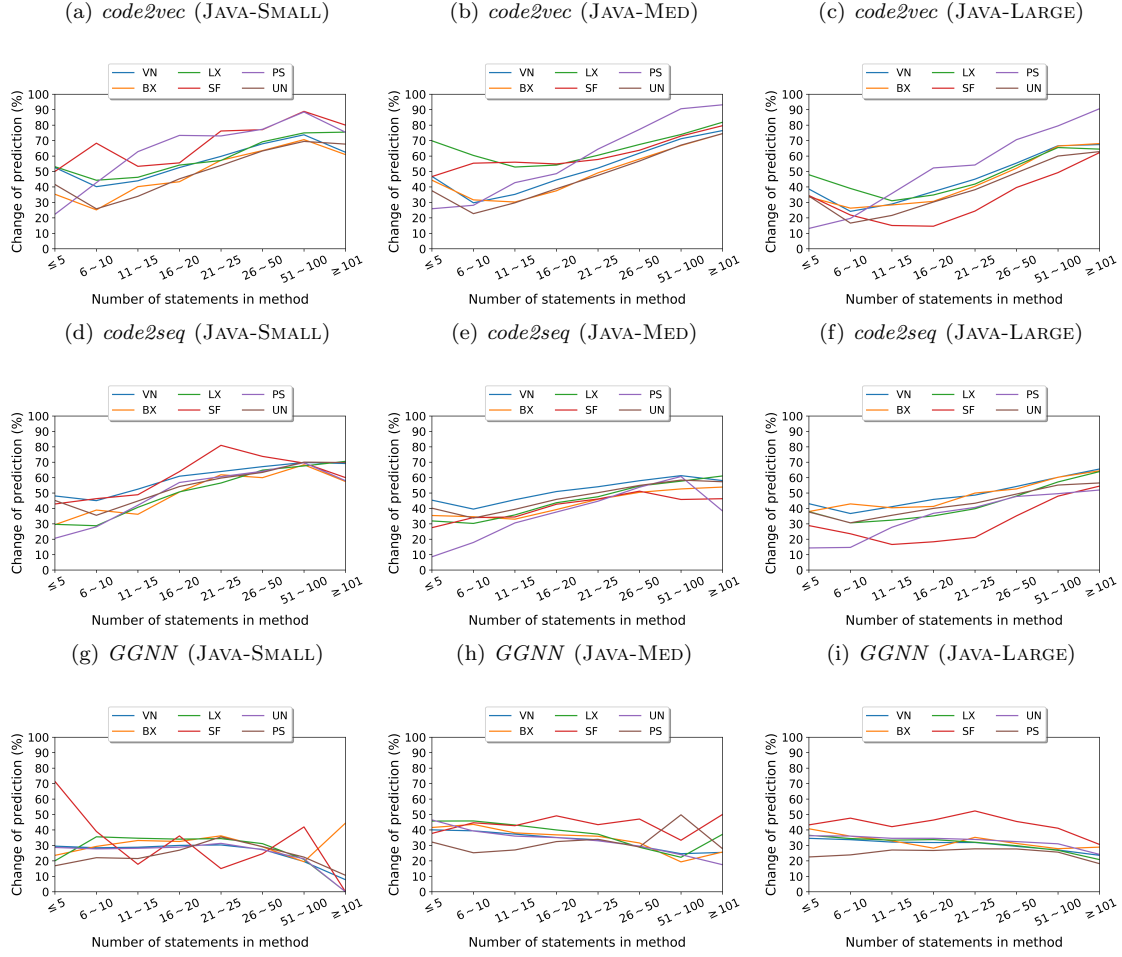


Figure 4: Prediction Change Percentage (PCP) across the number of statements in methods.

“Number of statements in method” denotes the number of executable lines in the body of methods before the transformation.

As shown in Figure 4(a-f), in most cases, the *code2vec* and *code2seq* models exhibit notable increases in PCP for all the transformations and datasets as the number of lines in methods increases. However, looking at Figure 4(g-i), it seems that *GGNN* is less sensitive to the number of lines in methods compared to *code2vec* and *code2seq* with respect to the transformations.

Observation 6: The *code2vec* and *code2seq* show notable increases in PCP as the length of methods grows, but PCP in *GGNN* seems to be less sensitive to the length of methods.

Table 4: The detailed PCP across all models, datasets, and transformations.

Dataset	Transformation	CCP			CWP			WWSP			WCP			WWDP		
		<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>
JAVA-SMALL	<i>Variable Renaming</i>	2.32	3.75	15.76	1.18	1.65	0.59	42.76	39.09	56.08	0.57	1.21	0.54	53.17	54.30	27.03
	<i>Boolean Exchange</i>	3.88	4.54	22.25	0.72	2.57	0.43	42.26	41.15	48.38	1.38	1.65	0.58	51.76	50.09	28.36
	<i>Loop Exchange</i>	1.86	4.24	16.23	0.74	0.74	0.89	38.76	43.22	52.11	0.72	1.10	0.62	57.92	50.70	30.15
	<i>Switch to If</i>	1.54	2.70	26.61	1.16	1.54	1.61	29.73	35.52	41.94	0.00	1.93	0.40	67.57	58.31	29.44
	<i>Permute Statement</i>	4.64	3.96	16.86	1.35	1.28	0.80	22.57	38.72	56.78	1.93	1.27	0.27	69.51	54.77	25.29
	<i>Unused Statement</i>	8.08	10.40	20.97	1.73	3.20	0.82	51.96	44.00	50.69	0.60	1.40	0.79	37.63	41.00	26.73
JAVA-MED	<i>Variable Renaming</i>	7.56	9.41	20.39	2.81	3.90	1.18	45.89	41.85	43.65	0.80	1.54	1.22	42.94	43.30	33.56
	<i>Boolean Exchange</i>	12.76	13.90	27.83	1.79	1.91	1.00	36.89	41.39	38.42	1.45	1.97	1.03	47.11	40.83	31.72
	<i>Loop Exchange</i>	6.61	7.62	20.93	2.21	1.71	1.09	30.62	47.09	42.40	1.22	1.33	1.23	59.34	42.25	34.35
	<i>Switch to If</i>	11.15	17.90	31.28	5.99	2.94	2.93	28.94	40.51	24.99	2.45	2.27	1.85	51.47	36.38	38.95
	<i>Permute Statement</i>	11.53	14.55	25.47	4.31	1.57	1.09	23.03	42.81	40.44	2.05	1.57	1.29	59.08	39.50	31.71
	<i>Unused Statement</i>	16.90	21.58	25.07	3.24	5.22	1.83	47.30	38.18	32.13	0.79	1.41	1.67	31.77	33.61	39.30
JAVA-LARGE	<i>Variable Renaming</i>	15.40	14.54	14.57	3.25	4.69	0.77	42.55	38.42	53.51	1.40	2.32	1.13	37.40	40.03	30.02
	<i>Boolean Exchange</i>	11.33	10.14	13.97	2.09	2.92	0.79	40.87	38.42	54.06	2.00	2.92	1.13	43.71	45.60	30.05
	<i>Loop Exchange</i>	19.81	18.87	13.29	3.00	2.49	0.28	33.67	38.62	54.96	2.13	2.30	1.13	41.39	37.72	30.34
	<i>Switch to If</i>	48.30	52.23	9.30	5.63	5.28	0.27	21.37	18.68	45.20	4.09	2.80	0.68	20.61	21.01	44.55
	<i>Permute Statement</i>	11.13	13.35	21.89	6.66	2.17	1.27	24.49	44.72	51.79	2.89	2.82	0.96	54.83	36.94	24.09
	<i>Unused Statement</i>	24.31	26.57	22.28	3.75	5.95	1.19	44.47	35.99	42.05	0.85	1.81	1.24	26.62	29.68	33.24

6.4. RQ4: Trends in the Types of Changes

Table 4 shows the full breakdown of the proportion of different types of changes after the transformation of methods. In this experiment, we use the same single-place transformed data that have been used for the PCP in Table 2. In *code2vec* and *code2seq*, the value of CCP increases with increase in the size of datasets. It may suggest that with a larger dataset the neural program model can generalize the correct predictions better.

In addition, we calculate $\frac{CWP}{CCP+CWP}$ to approximate the ratio of cases that the neural program model’s prediction switches from correct to wrong after transformations with respect to all the cases whose initial predictions are correct. The ratio helps us to simplify the comparison of (in)generalizability across different models. On average, 23% and 20% of cases, the neural program model switches from a correct prediction to a wrong one in *code2vec* and *code2seq*, respectively. In *GGNN*, on the other hand, this switch happens in less than 5% of transformations.

Similarly, $\frac{WCP}{WWSWP+WWDWP+WCP}$ approximates the ratio of cases switching from a wrong prediction to a correct prediction after transformations with respect to all the cases whose initial prediction are wrong. In *code2vec* and *code2seq*, a transformation switches from a wrong prediction to correct prediction in less than 3% of cases, however, this switch happens in around 1% of transformations for *GGNN*. Higher $\frac{CWP}{CCP+CWP}$ than $\frac{WCP}{WWSWP+WWDWP+WCP}$ implies that transformations are likely to reduce the overall performance of the neural program models.

Observation 7: Transformations are likely to decrease the overall performance of neural program models, and they are more likely to change the correct prediction in *code2vec* and *code2seq* than *GGNN*, while the generalizability of *code2vec* and *code2seq* can be compensated by larger datasets more than *GGNN*.

6.5. RQ5: Impact of the Transformations on Precision, Recall, and F_1 -Score

The performance of neural program models in the literature are often measured in classic metrics, such as precision, recall, and F_1 -score. In particular for the method name prediction task trained for *code2vec*, *code2seq* and *GGNN*, subtoken-level comparison is used to calculate the metrics; i.e., the method names in both predicted results and ground-truth names are split into individual tokens for the measurements (cf. the definitions in Section 4.3).

We also study the impact of the program transformations on the performance of neural program models in terms of these classic metrics. Table 5 shows the changed precision, recall, and F_1 -scores for the programs transformed by different transformations. In this

Table 5: The precision, recall and F_1 -score for subtokens across all models, datasets, and transformations.

Dataset	Transformation	# Transformed methods	Precision			Recall			F_1 -Score		
			<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>	<i>code2vec</i>	<i>code2seq</i>	<i>GGNN</i>
JAVA-SMALL	<i>Variable Renaming</i>	123123	9.79	38.01	40.64	5.05	28.99	23.78	6.66	32.89	30.00
	<i>Boolean Exchange</i>	1519	8.97	33.58	41.19	5.36	26.94	25.93	6.71	29.90	31.83
	<i>Loop Exchange</i>	5160	9.08	34.52	39.50	5.22	26.08	23.40	6.63	29.71	29.39
	<i>Switch to If</i>	259	7.01	30.78	38.32	4.99	26.41	30.16	5.83	28.43	33.75
	<i>Permute Statement</i>	9169	11.21	33.11	45.73	5.64	25.52	22.95	7.50	28.82	30.56
	<i>Unused Statement</i>	44426	21.26	50.99	44.37	13.63	41.12	31.36	16.61	45.53	36.75
	Weighted Average =		12.60	40.76	41.77	7.16	31.65	25.59	9.11	35.62	31.66
JAVA-MED	<i>Variable Renaming</i>	771208	20.90	43.57	39.71	10.89	28.98	22.24	14.32	34.81	28.51
	<i>Boolean Exchange</i>	8840	22.29	40.72	42.26	13.95	29.02	24.72	17.16	33.89	31.19
	<i>Loop Exchange</i>	23533	18.29	39.25	42.00	10.06	26.55	23.10	12.98	31.67	29.81
	<i>Switch to If</i>	3839	30.24	51.49	47.67	20.89	39.56	34.65	24.71	44.74	40.13
	<i>Permute Statement</i>	44711	24.29	38.75	39.79	12.99	28.26	23.35	16.93	32.68	29.43
	<i>Unused Statement</i>	351621	32.87	55.79	44.11	21.56	43.15	30.03	26.04	48.66	35.73
	Weighted Average =		24.51	46.88	41.09	14.12	33.08	24.63	17.87	38.74	30.74
JAVA-LARGE	<i>Variable Renaming</i>	916565	35.17	48.79	32.80	20.83	38.14	20.46	26.16	42.81	25.20
	<i>Boolean Exchange</i>	12107	27.30	42.90	23.82	15.83	33.19	16.54	20.04	37.43	19.52
	<i>Loop Exchange</i>	49665	37.60	46.79	25.28	23.91	37.91	18.75	29.23	41.88	21.53
	<i>Switch to If</i>	11165	69.34	72.75	22.68	57.06	66.18	21.81	62.60	69.31	22.24
	<i>Permute Statement</i>	74973	25.56	43.88	32.56	14.62	33.28	21.37	18.60	37.85	25.80
	<i>Unused Statement</i>	370927	44.96	61.57	45.43	30.40	52.44	29.93	36.27	56.64	36.09
	Weighted Average =		37.48	51.90	35.64	23.32	41.75	22.87	28.72	46.25	27.85

experiment, we use the same single-place transformed data that have been used for the PCP in Table 2.

In comparison with Table 1, we can see the average precision, recall, and F_1 -score in Table 5 have obvious decreases for all the three neural program models across the three Java datasets, that may indicate the (negative) impact of the transformations on the neural program models.

We also find no obvious correspondence between the PCP shown in Table 2 and the changes in precision, recall, and F_1 -score; high PCP does not necessarily lead to high changes in precision, recall and F_1 -score and vice versa.

Observation 8: Neural program models seem susceptible to semantic-preserving transformations with respect to the classic metrics of precision, recall, and F_1 -score as well. While our new metric of Prediction Change Percentage (PCP) shows the impact of the transformations from a different and more fine-grained perspective, the changes in the classic metrics are not correlated with PCP.

7. Discussion

In this paper, we study the current state of generalizability in neural program models built on *code2vec*, *code2seq*, and *GGNN*. Although limited, it provides interesting insights. In this section,

we discuss why neural networks have become a popular, or perhaps the de-facto, tool for processing programs, and what are the implications of using neural networks in processing source code.

Neural networks constitute a powerful class of machine learning models with a large hypothesis class. For instance, a multi-layer feed-forward network is called a universal approximator, meaning that it can essentially represent any function [25]. Unlike traditional learning techniques that require extensive feature engineering and tuning, deep neural networks facilitate representation learning. That is, they are capable of performing feature extraction out of raw data on their own [32]. Given a sufficiently large dataset, neural networks with adequate capabilities can substantially reduce the burden of feature engineering. Availability of a large number of code repositories makes data-driven program analysis a good application of neural networks. However, it is still unknown if neural networks are the best way to process programs [24] vs. [29].

Although the large hypothesis class of neural networks and feature learning make them very appealing to use, the complex models built by neural networks are still too difficult to understand and interpret. Therefore, as we apply neural networks in program analysis, we should develop specialized tools and techniques to enhance their interpretability, generalizability and robustness.

7.1. Generalizability vs. Interpretability vs. Robustness and Others

Interpretability studied in the literature may help to build more understandable neural networks, revealing the limits and strengths of the networks, and thus to some extent, it helps to evaluate and understand the generalizability of the networks. However, our study of generalizability with respect to program transformation provides a different perspective complement to interpretability; the approach may have the potential in the future to help identify interpretable code elements by measuring the impact of certain types of code transformations.

As mentioned in Section 2, there is a substantial line of work on evaluating the robustness of neural networks especially in the domain of vision and pattern recognition [46]. The key insight in such domains is that small, imperceptible changes in input should not impact the result of output. While this observation can be true for domains such as vision, it might not be directly applicable to the discrete domain of neural program models, since some minor changes to a program can drastically change the semantic and behavior of the program. Quantifying the imperceptibility and many other aspects of source code is our future research goal.

7.2. Are we there yet?

Are neural program models ready for widespread use in program analysis? The neural program models in our experiments are brittle to even very small changes in the methods. The semantic-preserving transformations can change the outputs of the neural program models in 26% to 73% of cases. Although our findings are limited to only one task, they suggest caution. The literature lacks techniques for rigorous evaluation of neural program models. The recent line of work by Nghi et al. [13] in interpretability of neural program models, Rabin et al. [40] in testing them, and Yefet et al. [56] are much needed steps in a right direction.

7.3. Code Representation

The performance of models used in neural program models, such as ones used in this study, is relatively low compared to the performance of neural models in domains such as natural language understanding [42], text classification [31]. To improve their performance, we would need novel code representations that better capture interesting characteristics of programs.

8. Related Work

Robustness of Neural Networks. There is a substantial line of work on the robustness of artificial intelligence (AI) systems in general and deep neural networks in particular. Szegedy et al. [46] is the first to discover that deep neural networks are vulnerable to small perturbations that are imperceptible to human eyes. They developed the L-BFGS method for the systematic generation of such adversarial examples. Goodfellow et al. [21] proposes a more efficient method, called the Fast Gradient Sign Method that exploits the linearity of deep neural networks. **Many following up works [30, 37, 14, 18, 57, 61, 56] further demonstrated the severity of the robustness issues with a variety of attacking methods and defenses.** While aforementioned approaches only apply to models for image classification, new attacks have been proposed that target models in other domains, such as natural language processing [33, 27, 62] and graphs [16, 63].

The automated verification research community has proposed techniques to offer guarantees for the robustness of neural networks by adapting bounded model checking [44], abstract interpretation [20], and Satisfiability Modulo Theory [26]. Amershi et al. [8] study the challenges in developing AI solutions and Zhang et al. [59] survey testing of machine-learning systems.

Models of Code. Early works directly adopted NLP models to discover textual patterns existed in the source code [22, 39]. Those methods, unfortunately, do not account for the structural

information programs exhibit. Following approaches address this issue by generalizing from the abstract syntax trees [36, 38, 7, 6]. As Graph Neural Networks (GNN) [43] have been gaining increasing popularity due to its remarkable representation capacity, many works have leveraged GNN to tackle challenging tasks like program repair and bug finding, and obtained quite promising results [4, 53, 17, 34, 12, 17, 47]. Besides, the attention mechanism [9] has been applied into GNNs to improve the performance further [10, 54, 60]. It is very interesting to see how the attention can help to explain the output of the neural models [48, 55, 13]. In parallel, Wang *et al.* developed a number of models [49, 51, 52] that feed off the run time information for enhancing the precision of semantic representation for model inputs.

9. Threats to Validity

There are various threats to the validity of our approach.

Limited Data and Evaluation Scope. We only evaluated the generalizability of neural program models built on *code2vec*, *code2seq*, and *GGNN*, for one task in Java programs. Therefore, our results may not generalize to other neural program models or other tasks or other programming languages. We leave the evaluation of the general applicability of our approach as future work.

Transformations. The proposed transformations in this paper impact program ASTs in varying degrees. Some of the transformations, e.g. variable renaming, are common refactoring techniques. However, these transformations may not represent many possible transformations in other domains. We will instantiate and extend our approach with other transformations from other domains.

Internal Validity. Some bugs may exist in the toolchain and neural program models implemented in this paper. To reduce the probability of bugs, two authors reviewed the code and manually inspected a sample of transformed programs to ensure the reliability of transformations.

10. Conclusion & Future Work

In this paper, we perform a large-scale, systematic evaluation of the generalizability of state-of-the-art neural program models built on *code2vec*, *code2seq*, and *GGNN*. In particular, we apply six semantic-preserving program transformations to produce new programs on which we expect the neural program models to keep their original predictions. We find that such program transformations frequently sway the predictions of these neural program models, indicating serious generalization

issues that could negatively impact the wider applications of deep neural networks in program analysis tasks. Although neural program models that encode more program dependency information and are trained with larger datasets may exhibit more generalizable behavior, their generalizability is still limited. We believe this work provides a systematic approach and metrics for evaluating neural program models, and can motivate future research on training not only accurate but also generalizable deep models of code. Future work that includes more semantic-preserving and even some semi-semantic-preserving transformations in our approach and adapts more fine-grained prediction change metrics may further extend the applicability of our approach to various neural program models designed for different tasks. We also plan to explore using transformed programs to improve the generalizability of the neural program models.

References

- [1] Allamanis, M., 2019. The adverse effects of code duplication in machine learning models of code, in: Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software, Association for Computing Machinery, New York, NY, USA. p. 143–153. URL: <https://doi.org/10.1145/3359591.3359735>, doi:10.1145/3359591.3359735.
- [2] Allamanis, M., Barr, E.T., Bird, C., Sutton, C., 2015. Suggesting accurate method and class names, in: Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, Association for Computing Machinery, New York, NY, USA. p. 38–49. URL: <https://doi.org/10.1145/2786805.2786849>, doi:10.1145/2786805.2786849.
- [3] Allamanis, M., Barr, E.T., Devanbu, P., Sutton, C., 2018a. A survey of machine learning for big code and naturalness. *ACM Comput. Surv.* 51. URL: <https://doi.org/10.1145/3212695>, doi:10.1145/3212695.
- [4] Allamanis, M., Brockschmidt, M., Khademi, M., 2018b. Learning to represent programs with graphs, in: International Conference on Learning Representations.
- [5] Allamanis, M., Peng, H., Sutton, C.A., 2016. A convolutional attention network for extreme summarization of source code, in: Proceedings of the 33rd International Conference on Machine Learning, ICML.

- [6] Alon, U., Levy, O., Yahav, E., 2019a. code2seq: Generating sequences from structured representations of code, in: International Conference on Learning Representations.
- [7] Alon, U., Zilberstein, M., Levy, O., Yahav, E., 2019b. Code2vec: Learning distributed representations of code. Proc. ACM Program. Lang. 3, 40:1–40:29. URL: <http://doi.acm.org/10.1145/3290353>, doi:10.1145/3290353.
- [8] Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., Zimmermann, T., 2019. Software engineering for machine learning: A case study, in: Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice, IEEE Press. p. 291–300. URL: <https://doi.org/10.1109/ICSE-SEIP.2019.00042>, doi:10.1109/ICSE-SEIP.2019.00042.
- [9] Bahdanau, D., Cho, K., Bengio, Y., 2015. Neural machine translation by jointly learning to align and translate, in: Bengio, Y., LeCun, Y. (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.
- [10] Beck, D., Haffari, G., Cohn, T., 2018. Graph-to-sequence learning using gated graph neural networks, in: Proceedings of the Association for Computational Linguistics (ACL), Association for Computational Linguistics. pp. 273–283.
- [11] Bielik, P., Vechev, M., 2020. Adversarial robustness for code. Proceedings of the International Conference on Machine Learning (ICML) .
- [12] Brockschmidt, M., Allamanis, M., Gaunt, A.L., Polozov, O., 2019. Generative code modeling with graphs, in: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019.
- [13] Bui, N.D.Q., Yu, Y., Jiang, L., 2019. Autofocus: Interpreting attention-based neural networks by code perturbation, in: 2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 38–41. doi:10.1109/ASE.2019.00014.
- [14] Carlini, N., Wagner, D., 2017. Towards evaluating the robustness of neural networks, in: 2017 IEEE Symposium on Security and Privacy (SP), IEEE. pp. 39–57.

- [15] Compton, R., Frank, E., Patros, P., Koay, A., 2020. Embedding java classes with code2vec: Improvements from variable obfuscation, in: 17th International Conference on Mining Software Repositories (MSR) 2020, Seoul, Republic of Korea.
- [16] Dai, H., Li, H., Tian, T., Huang, X., Wang, L., Zhu, J., Song, L., 2018. Adversarial attack on graph structured data. Proceedings of the 35th International Conference on Machine Learning, PMLR 80.
- [17] Dinella, E., Dai, H., Li, Z., Naik, M., Song, L., Wang, K., 2019. Hoppity: Learning graph transformations to detect and fix bugs in programs, in: International Conference on Learning Representations.
- [18] Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J., 2018. Boosting adversarial attacks with momentum, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 9185–9193.
- [19] Fernandes, P., Allamanis, M., Brockschmidt, M., 2019. Structured neural summarization, in: International Conference on Learning Representations.
- [20] Gehr, T., Mirman, M., Drachler-Cohen, D., Tsankov, P., Chaudhuri, S., Vechev, M., 2018. Ai2: Safety and robustness certification of neural networks with abstract interpretation, in: 2018 IEEE Symposium on Security and Privacy (SP), IEEE. pp. 3–18.
- [21] Goodfellow, I.J., Shlens, J., Szegedy, C., 2015. Explaining and harnessing adversarial examples, in: Bengio, Y., LeCun, Y. (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.
- [22] Gupta, R., Pal, S., Kanade, A., Shevade, S., 2017. Deepfix: Fixing common c language errors by deep learning, in: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI Press. p. 1345–1351.
- [23] Hellendoorn, V.J., Bird, C., Barr, E.T., Allamanis, M., 2018. Deep learning type inference, in: Proceedings of the 2018 26th acm joint meeting on european software engineering conference and symposium on the foundations of software engineering, pp. 152–162.
- [24] Hellendoorn, V.J., Devanbu, P., 2017. Are deep neural networks the best choice for modeling source code?, in: Proceedings of the 2017 11th Joint Meeting on Foundations of Software

Engineering, Association for Computing Machinery, New York, NY, USA. p. 763–773. URL: <https://doi.org/10.1145/3106237.3106290>, doi:10.1145/3106237.3106290.

- [25] Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2, 359–366.
- [26] Huang, X., Kwiatkowska, M., Wang, S., Wu, M., 2017. Safety verification of deep neural networks, in: Majumdar, R., Kunčák, V. (Eds.), *Computer Aided Verification*, Springer International Publishing, Cham. pp. 3–29.
- [27] Jia, R., Liang, P., 2017. Adversarial examples for evaluating reading comprehension systems, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. URL: <https://doi.org/10.18653/v1/d17-1215>.
- [28] Kang, H.J., Bissyandé, T.F., Lo, D., 2019. Assessing the generalizability of code2vec token embeddings, in: *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pp. 1–12. doi:10.1109/ASE.2019.00011.
- [29] Karampatsis, R.M., Sutton, C., 2019. Maybe deep neural networks are the best choice for modeling source code. *arXiv:1903.05734*.
- [30] Kurakin, A., Goodfellow, I.J., Bengio, S., 2017. Adversarial machine learning at scale, in: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- [31] Lai, S., Xu, L., Liu, K., Zhao, J., 2015. Recurrent convolutional neural networks for text classification, in: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI Press. p. 2267–2273.
- [32] LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *nature* 521, 436.
- [33] Li, J., Monroe, W., Jurafsky, D., 2016a. Understanding neural networks through representation erasure. *arXiv preprint arXiv:1612.08220* .
- [34] Li, Y., Tarlow, D., Brockschmidt, M., Zemel, R.S., 2016b. Gated graph sequence neural networks, in: Bengio, Y., LeCun, Y. (Eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

- [35] Liu, K., Kim, D., Bissyandé, T.F., Kim, T., Kim, K., Koyuncu, A., Kim, S., Traon, Y.L., 2019. Learning to spot and refactor inconsistent method names, in: Proceedings of the 41st International Conference on Software Engineering, IEEE Press. p. 1–12. URL: <https://doi.org/10.1109/ICSE.2019.00019>, doi:10.1109/ICSE.2019.00019.
- [36] Maddison, C., Tarlow, D., 2014. Structured generative models of natural source code, in: International Conference on Machine Learning, pp. 649–657.
- [37] Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P., 2016. Deepfool: a simple and accurate method to fool deep neural networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2574–2582.
- [38] Mou, L., Li, G., Zhang, L., Wang, T., Jin, Z., 2016. Convolutional neural networks over tree structures for programming language processing, in: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI Press. p. 1287–1293.
- [39] Pu, Y., Narasimhan, K., Solar-Lezama, A., Barzilay, R., 2016. sk_p: a neural program corrector for moocs, in: Companion Proceedings of the 2016 ACM SIGPLAN International Conference on Systems, Programming, Languages and Applications: Software for Humanity, ACM.
- [40] Rabin, M.R.I., Wang, K., Alipour, M.A., 2019. Testing neural program analyzers. 34th IEEE/ACM International Conference on Automated Software Engineering (Late Breaking Results-Track) URL: <https://arxiv.org/abs/1908.10711>.
- [41] Ramakrishnan, G., Henkel, J., Wang, Z., Albarghouthi, A., Jha, S., Reps, T., 2020. Semantic robustness of models of source code. arXiv preprint arXiv:2002.03043 .
- [42] Sarikaya, R., Hinton, G.E., Deoras, A., 2014. Application of deep belief networks for natural language understanding. IEEE/ACM Trans. Audio, Speech and Lang. Proc. 22, 778–784. URL: <https://doi.org/10.1109/TASLP.2014.2303296>, doi:10.1109/TASLP.2014.2303296.
- [43] Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G., 2009. The graph neural network model. IEEE Transactions on Neural Networks 20, 61–80. doi:10.1109/TNN.2008.2005605.
- [44] Scheibler, K., Winterer, L., Wimmer, R., Becker, B., 2015. Towards verification of artificial neural networks., in: MBMV, pp. 30–40.

- [45] Sundararajan, M., Taly, A., Yan, Q., 2017. Axiomatic attribution for deep networks, in: Proceedings of the 34th International Conference on Machine Learning, pp. 3319–3328.
- [46] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2014. Intriguing properties of neural networks, in: ICLR.
- [47] Tarlow, D., Moitra, S., Rice, A., Chen, Z., Manzagol, P.A., Sutton, C., Aftandilian, E., 2020. Learning to fix build errors with graph2diff neural networks, in: Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops, pp. 19–20.
- [48] Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y., 2018. Graph attention networks, in: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings.
- [49] Wang, K., 2019. Learning scalable and precise representation of program semantics. arXiv preprint arXiv:1905.05251 .
- [50] Wang, K., Christodorescu, M., 2019. Coset: A benchmark for evaluating neural program embeddings. arXiv:1905.11445.
- [51] Wang, K., Singh, R., Su, Z., 2018. Dynamic neural program embedding for program repair, in: International Conference on Learning Representations.
- [52] Wang, K., Su, Z., 2020. Blended, precise semantic program embeddings, in: Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation.
- [53] Wang, Y., Gao, F., Wang, L., Wang, K., 2019. Learning a static bug finder from data. arXiv preprint arXiv:1907.05579 .
- [54] Xu, K., Wu, L., Wang, Z., Feng, Y., Witbrock, M., Sheinin, V., 2018a. Graph2seq: Graph to sequence learning with attention-based neural networks. arXiv preprint arXiv:1804.00823 .
- [55] Xu, X., Zu, S., Gao, C., Zhang, Y., Feng, W., 2018b. Modeling attention flow on graphs. Relational Representation Learning Workshop, NeurIPS 2018 .

- [56] Yefet, N., Alon, U., Yahav, E., 2020. Adversarial examples for models of code. Proc. ACM Program. Lang. 4. URL: <https://doi-org.ezproxy.lib.uh.edu/10.1145/3428230>, doi:10.1145/3428230.
- [57] Yuan, X., He, P., Zhu, Q., Li, X., 2019. Adversarial examples: Attacks and defenses for deep learning. IEEE transactions on neural networks and learning systems 30, 2805–2824.
- [58] Zhang, H., Li, Z., Li, G., Ma, L., Liu, Y., Jin, Z., 2020. Generating adversarial examples for holding robustness of source code processing models, in: The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI Press. pp. 1169–1176.
- [59] Zhang, J.M., Harman, M., Ma, L., Liu, Y., 2020. Machine learning testing: Survey, landscapes and horizons. IEEE Transactions on Software Engineering , 1–doi:10.1109/TSE.2019.2962027.
- [60] Zhang, S., Xie, L., 2020. Improving attention mechanism in graph neural networks via cardinality preservation, in: Bessiere, C. (Ed.), Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, ijcai.org. pp. 1395–1402. URL: <https://doi.org/10.24963/ijcai.2020/194>, doi:10.24963/ijcai.2020/194.
- [61] Zhang, X., Zitnik, M., 2020. Gnn-guard: Defending graph neural networks against adversarial attacks. Advances in Neural Information Processing Systems 33.
- [62] Zhao, Z., Dua, D., Singh, S., 2018. Generating natural adversarial examples, in: International Conference on Learning Representations.
- [63] Zügner, D., Akbarnejad, A., Günnemann, S., 2018. Adversarial attacks on neural networks for graph data, in: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2847–2856.