funcGNN: A Graph Neural Network Approach to Program Similarity

Aravind Nair

KTH Royal Institute of Technology
Stockholm, Sweden
aanair@kth.se

Avijit Roy

KTH Royal Institute of Technology
Stockholm, Sweden
avijit@kth.se

Karl Meinke

KTH Royal Institute of Technology
Stockholm, Sweden
karlm@kth.se

ABSTRACT

Background: Program similarity is a fundamental concept, central to the solution of software engineering tasks such as software plagiarism, clone identification, code refactoring and code search. Accurate similarity estimation between programs requires an in-depth understanding of their structure, semantics and flow. A control flow graph (CFG), is a graphical representation of a program which captures its logical control flow and hence its semantics. A common approach is to estimate program similarity by analysing CFGs using graph similarity measures, e.g. graph edit distance (GED). However, graph edit distance is an NP-hard problem and computationally expensive, making the application of graph similarity techniques to complex software programs impractical. Aim: This study intends to examine the effectiveness of graph neural networks to estimate program similarity, by analysing the associated control flow graphs. Method: We introduce *funcGNN*¹, which is a graph neural network trained on labeled CFG pairs to predict the GED between unseen program pairs by utilizing an effective embedding vector. To our knowledge, this is the first time graph neural networks have been applied on labeled CFGs for estimating the similarity between high-level language programs. Results: We demonstrate the effectiveness of *funcGNN* to estimate the GED between programs and our experimental analysis demonstrates how it achieves a lower error rate (1.94×10^{-3}), with faster (23 times faster than the quickest traditional GED approximation method) and better scalability compared with state of the art methods. Conclusion: *funcGNN* possesses the inductive learning ability to infer program structure and generalise to unseen programs. The graph embedding of a program proposed by our methodology could be applied to several related software engineering problems (such as code plagiarism and clone identification) thus opening multiple research directions.

CCS CONCEPTS

• Information systems → Similarity measures; • Computing methodologies → Neural networks; Semantic networks.

¹The code is publicly available at <https://github.com/aravi11/funcGNN>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEM '20, October 8–9, 2020, Bari, Italy

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7580-1/20/10...\$15.00

<https://doi.org/10.1145/3382494.3410675>

KEYWORDS

Software Engineering, Machine Learning, Program Similarity, Graph Neural Network, Graph Similarity, Graph Edit Distance, Attention Mechanism, Graph Embedding, Control Flow Graph

ACM Reference Format:

Aravind Nair, Avijit Roy, and Karl Meinke. 2020. funcGNN: A Graph Neural Network Approach to Program Similarity. In *ESEM '20: ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM) (ESEM '20), October 8–9, 2020, Bari, Italy*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3382494.3410675>

1 INTRODUCTION

Finding the similarity between two objects plays an important role in many computational tasks such as recommendation and marketing. For example, well-known search engines (Google, Bing), e-commerce sites (Amazon, e-Bay) and online media service providers (YouTube, Netflix), all invest heavily in similarity measures to recommend the next best product for their users. Another area where similarity measures are helpful is in software engineering. Program similarity or code similarity is a fundamental theoretical concept, central to the effective solution of software engineering tasks such as software plagiarism, clone identification, code refactoring and code search. The basic idea of a program similarity metric is to quantitatively measure how one program is syntactically similar to another program. Program similarity is related to program equivalence, where the latter concept usually refers to semantic similarity. However, program equivalence is a more challenging task as a minor change in code structure can drastically change its logic making it structurally similar but semantically very different. Since semantic equivalence of programs is undecidable, program similarity represents a simpler but more tractable syntactic approximation.

One of the most widely used techniques to analyse programs is by transforming them into graphs [1]. A graph is a mathematical structure used to represent the relationship and connection between objects termed nodes. Graph structures are ubiquitous in real-life and can be found in almost every domain. The most frequently used graph representations in the field of program analysis are control graphs (CG), abstract syntax trees (AST), control flow graphs (CFG) and program dependency graphs (PDG). A CFG is a directed graph in which each node represents an atomic operation or statement, and each edge represents a possible transfer of control (i.e. execution order). As CFGs are capable of preserving both the logic flow and semantics of a program, we can use the CFG representation to address the problem of program similarity. We have used Soot² as an open source bytecode manipulation

²<https://github.com/Sable/soot>

Figure 1: Java function to calculate sum of all elements in an array and its corresponding CFG

and optimization framework to generate CFGs for Java programs. Soot processes the bytecode of a Java program and converts it into an intermediate representation *Jimple* [47]. *Jimple* breaks down a statement to a 3-address instruction set to provide a detailed atomic operation model of the program. These atomic operations label the nodes of the CFG. Figure 1 shows the transformation of a simple Java function (sum of all elements in an array) to its corresponding CFG.

Though graphs provide good program models, graph similarity is not an easy task. Graph similarity metrics like Graph Edit Distance (GED) [6] and Maximum Common Subgraph (MCS) [7] are known to be NP-hard problems, hence computationally expensive. Using the GED metric for large graphs is impractical, as one cannot compute the similarity score between graphs of more than 16 nodes within a reasonable time [4]. Over the past ten years, graph neural networks (GNN) have emerged in machine learning (ML) as a successful class of neural network models. GNNs are capable of supervised learning of graph representations to efficiently implement many relations and functions on graphs. In one approach, a GNN generates a meaningful vector embedding for each node in a graph using a recursive neighborhood aggregation algorithm [45]. This embedding approximates the semantics of the corresponding nodes and can be used for supervised machine learning tasks like node classification, graph classification, graph similarity etc.

In this research, we aim to investigate whether program similarity could be estimated by analysing labeled control flow graphs using a graph neural network. We introduce *funcGNN*, a graph neural network trained to predict the GED between program pairs. In our approach, we use an amalgamation of two graph embedding techniques. The first technique is a *top-down* approach which creates an embedding for the whole graph by transforming into a meaningful fixed size vector. To make sure that this embedding

vector captures the semantic information of semantically significant nodes we incorporate an attention mechanism into the neural network architecture. The second technique, on the other hand, uses a *bottom-up* approach. This focuses on node level comparison of graph pairs and captures the atomic program operation similarities. Using these atomic node-level similarities, we compute a histogram feature representation which gives an inferred probability distribution for node similarity. The two vector embeddings generated by these two techniques are concatenated, and the resulting embedding is the input vector for multiple fully connected neural network layers to predict the similarity score between a pair of graphs. The similarity score thus predicted is the normalized GED score between the graph pair.

To our knowledge, this is the first time graph neural networks have been applied on labeled CFGs of high level languages for program similarity. We evaluate the effectiveness of our proposed methodology on functions from open source Java code.

The main contributions of this research can be summarized as follows:

- We address the problem of program similarity and propose *funcGNN*, a novel graph neural network capable of predicting the similarity between program pairs by analysing their labeled control flow graphs.
- We use two graph embedding techniques with an attention mechanism and histogram feature representation to calculate the overall program similarity.
- We empirically demonstrate how *funcGNN*: (i) generalises well to unseen graph program pairs, (ii) achieves low error rates and (iii) significantly reduces computation times compared to the state of the art GED approximation methods. Thus our solution has better scaling properties.

The rest of this paper is organised as follows. We briefly explain the theoretical prerequisites for our work in Section 2. This includes a brief introduction to control flow graphs, graph edit distance and graph neural networks. Section 3 describes in detail *funcGNN*, our proposed ML architecture along with its sub-modules and their algorithms. In Section 4, we describe our evaluation dataset for *funcGNN*, and the hyper parameters of *funcGNN*. We present the experimental results obtained on our dataset. We discuss the limitations and threats to validity of our study in Section 5. Section 6 provides a survey of relevant literature. We conclude the paper in Section 7 with a summary of our results and a discussion of future directions for research.

2 BACKGROUND

2.1 Control Flow Graphs

To estimate the similarity between two programs, we will compare their control flow graphs³. A CFG is a graph representation which specifies the logic and control flow of a program. A control flow graph can be represented by a directed graph $G = (V, E \subseteq V \times V)$, where V denotes the set of nodes and E the set of edges (node pairs) connecting them. Each node $v_i \in V$, is labelled by an atomic program statement. A pair of nodes v_i, v_j is connected by an edge

³https://en.wikipedia.org/wiki/Control-flow_graph

$(v_i, v_j) \in E$, when this reflects the direct execution order: v_i is immediately followed by v_j . Thus a CFG provides a graph representation of both program syntax and semantics⁴.

2.2 Graph Edit Distance and its Approximations

Though graphs are ubiquitous in almost every field of computer science, finding similar graphs is a challenging task. There are well defined graph similarity metrics like graph edit distance (GED) [6] and maximum common subgraph (MCS) [7]. In this research we use GED as the similarity metric between two graphs. Graph edit distance is analogous to the edit distance concept used for string matching [9, 44]. The GED of two graphs can be defined as the number of operations required to transform one graph to another. Formally, given two graphs G_1 and G_2 , then the GED between them can be defined by,

$$GED_{(G_1, G_2)} = \min_{(x_1, \dots, x_k) \in P(G_1, G_2)} \sum_{i=1}^k c(x_i) \quad (1)$$

where, $P(G_1, G_2)$ denotes the set of all possible edit paths (sequences of atomic edit operations) that transform G_1 to G_2 . Here, $c(x) \geq 0$ denotes the cost of each graph edit operation x , which includes deletion, insertion and substitution of nodes and edges. A classical approach to calculate the GED between two graphs is described in [38], in which they identify the minimal edit path using the A* algorithm. However this approach has exponential time complexity and several studies have been made to reduce its execution time [5, 12, 37, 42, 43, 53].

In this study, we estimate the approximate GED between graph pairs by using the Quadratic Assignment Problem (QAP) approximation as proposed in [5]. The inspiration to substitute approximate values in place of the exact GED for large graphs comes from the ICPR2016⁵ contest. Once the approximate graph edit distance GED between a graph pair is calculated, we normalize it in a range between 0 and 1. This normalized score acts as the similarity score or metric for that graph pair. A detailed explanation of the normalization function is provided in Section 4.1

2.3 Graph Neural Networks

The graph neural network (GNN) model was designed as an extension of the recurrent neural network model. The goal was to efficiently learn relations and functions on graphs [16, 45]. The concept of GNN is based on the idea that each node in a graph $G = (V, E)$ can be characterised in terms of: (i) its own features, (ii) the relations it has with its locally neighbouring nodes, and (iii) the features of its local neighbours. To represent each node $v \in V$, a GNN uses an s -dimensional state embedding vector $h_v \in \mathbb{R}^s$ which consists of information about the node and its neighbours. The state embedding h_v is learned via a parametric non-linear local transition function f_{lt} , which is uniformly defined parametrically across all nodes. This local transition function f_{lt} captures the above three characteristics of a node and its local neighbors. The state

embedding h_v can be mathematically defined as,

$$h_v = f_{lt}(x_v, x_{co[v]}, x_{ne[v]}, h_{ne[v]}) \quad (2)$$

where, x_v represents the features of node v , $x_{co[v]}$ represents the features of the edges of node v , $x_{ne[v]}$ represents the features of neighboring nodes of v and $h_{ne[v]}$ represents the states of the neighboring nodes of v . The state embedding h_v along with the node feature x_v is used to learn the final representation o_v of the node v , using a parametric non-linear local output function f_{lo} . This final representation o_v is defined as follows:

$$o_v = f_{lo}(h_v, x_v) \quad (3)$$

Let H, O, X and X_n be the tensor vectors generated by stacking the state, output, all features and individual node features $v \in V$ in the graph G . Then equations 2 & 3 can be written as:

$$H = F_{GT}(H, X) \quad (4)$$

$$O = F_{GO}(H, X_n) \quad (5)$$

where, F_{GT} representing the global transition function, and F_{GO} representing the global output function are the stacked versions of f_{lt} and f_{lo} respectively. A unique solution to equations 4 & 5 can be found by using Banach's Fixed Point Theorem [25]. According to this theorem, a unique solution can be calculated as a fixed point of the operators F_{GT} and F_{GO} provided that these are *contraction maps* with respect to the state. The contraction condition means that there exists some $0 \leq \mu < 1$ such that,

$$\|F_{GT}(H, X) - F_{GT}(I, X)\| \leq \mu \|H - I\| \quad (6)$$

for any H, I , where $\|\cdot\|$ denotes the vector norm on states. Besides guaranteeing a unique solution, Banach's Theorem actually gives an iterative scheme for approximating the fixed-point:

$$H^{t+1} = F_{GT}(H^t, X) \quad (7)$$

where, H^t denotes the t^{th} iteration of H . For any initial value of $H(0)$, the dynamical system of equation 7 converges exponentially fast to a fixed point solution of equation 4. Thus, $H(t)$ denotes the state that is updated by the global transition function F_{GT} . Equation 7 could then be rewritten as,

$$h_v^{t+1} = f_{lt}(x_v, x_{co[v]}, x_{ne[v]}, h_{ne[v]}^t) \quad (8)$$

$$o_v^t = f_{lo}(h_v^t, x_v) \quad (9)$$

By keeping the target information t_v for each node v , GNNs learn the parameters of f_{lt} and f_{lo} by minimising the loss between the targeted value t_v and the output value o_v . This loss function can be represented as,

$$loss = \sum_{i=1}^N (t_v - o_v)^2 \quad (10)$$

where N is the number of nodes in the graph G . The learning task is carried out iteratively based on a gradient-descent strategy until time T , where the fixed point solution of equation 4 is achieved i.e. $H(T) \approx H$.

⁴A CFG can be compared with an abstract syntax tree which represents only program syntax.

⁵<https://gdc2016.greyc.fr/>

3 PROPOSED METHODOLOGY

In this section, we first describe the *funcGNN* framework to learn the approximate GED between program pairs. Given two graphs G_1 and G_2 , *funcGNN* creates a fixed size vector embedding for each graph, and learns a model of the similarity function, which could map these input embeddings to a single real-valued similarity score. To generate the embeddings, we use a combination of the whole graph embedding (a top-down approach) and atomic node level comparison representation (a bottom-up approach). The embeddings generated by these two methods are concatenated and fed to multiple fully connected neural network layers to calculate the similarity score. Figure 2 depicts the overall architecture and workflow of *funcGNN*. A detailed explanation of this architecture is given in the sequel.

3.1 Top-down Approach - Overall Graph Embedding

In the top-down approach, we aim to efficiently create a global embedding for the entire CFG capturing its global structure and control flow. This global embedding of a CFG is then used to predict program similarity. This global approach involves the following stages:

3.1.1 Inductive Node Embedding: For generating the embedding of each node v in the CFG, we use the GraphSAGE method as proposed in [17]. GraphSAGE is an inductive node embedding method which generalizes to unseen nodes, which is one of the major problems when analysing a new program not previously seen in the training data. The GraphSAGE methodology is different from the original GNN approach [45] as it defines the neighborhood of a node by aggregating the features from a sampled subset of its entire neighborhood. This could be represented as,

$$h_{N_v}^t = AGG_t(h_u^{t-1}), \forall u \in N_v \quad (11)$$

$$h_v^t = \sigma(W_t \cdot [h_v^{t-1} || h_{N_v}^t]) \quad (12)$$

where N_v represents the neighborhood set of node v , AGG_t represents the aggregate function, W represents the weight matrix, $||$ represents the concatenation operation, σ represents the non-linear activation function and h_v^t denotes the state embedding of node v at time t . We use the *mean aggregator* as the aggregator function AGG_t which is an approximation of the convolution operation in the GCN framework proposed in [27]. The *mean aggregator* function is a variant of the skip connection [18] and does not perform the concatenation operation as in equation 12. Thus equation 12 could be rewritten as,

$$h_v^t = \sigma(W \cdot \text{MEAN}(\{h_v^{t-1}\} \cup \{h_u^{t-1}, \forall u \in N_v\})) \quad (13)$$

In our approach, the nodes are initially represented by the one-hot embedding scheme $O \in \mathbb{R}^{N \times D}$ where N is the number of nodes in the graph and D is the dimension size. This one-hot embedding is then passed through multiple layers of GraphSAGE to get their node representation $U \in \mathbb{R}^{N \times D}$. We set the GraphSAGE layers to 3 as GNNs have an over-smoothing issue with deep architectures [30] and we use *relu* as our activation function.

3.1.2 Attention Based Graph Embedding. Once we obtain the representations of each node U , the next task is to combine them effectively to generate an overall embedding h for the whole graph. Instead of simply averaging the embedding of all the nodes in the graph we use an attention mechanism to provide more significance to certain nodes based on a similarity metric. This approach makes sure that nodes with more structural significance have more influence on the overall graph embedding compared to the other nodes. From a program analysis point of view, the idea is to provide more weight to nodes labelled with mathematical operations than nodes labelled with variable initialization or assignment.

To achieve this, we first compute a context vector embedding c by averaging all the node embeddings U and transforming them through a non-linear *relu* activation function. This graph context vector c can be represented as,

$$c = \text{relu}\left(\frac{1}{|V|} \sum_{v=1}^{|V|} U_v \cdot W\right) \quad (14)$$

where $|V|$ is the number of nodes in the graph and W is a weight matrix of dimension $W \in \mathbb{R}^{D \times D}$. By learning the weight matrix W , the context vector c provides a naive summary of the structural attributes of the entire graph. To calculate the attention weight a_v of each node v we take the inner product of each node embedding U_v with the context vector c . The idea behind this approach is that majority of the node operations in a program will have mathematical operations rather than variable initialization or assignment operations. Hence the node embeddings for mathematical operations will have more impact on the context vector c , and node embeddings most similar to the context vector c will attain higher attention weight. Once we calculate the attention weight a_v for each node, we calculate the overall graph embedding $h \in \mathbb{R}^D$ by,

$$a_v = \text{relu}(U_v^T \cdot c) \quad (15)$$

$$h = \sum_{v=1}^{|V|} (a_v \cdot U_v) \quad (16)$$

3.1.3 NTN based Graph Pair Comparison: Now that we have obtained the overall embedding h for each graph, the challenge is to efficiently compare two graph embeddings to estimate their similarity. For this we use the *Neural Tensor Network (NTN)* as proposed in [46]. The advantage of the *neural tensor network (NTN)* over the traditional linear layer approach is its ability to efficiently compare two embedding vectors across multiple dimensions. Given two embedding vectors h_i and h_j , NTN makes use of a bilinear tensor, which computes the relationship between the two embeddings using the following function,

$$\text{NTN}(h_i, h_j) = \sigma(h_i^T \cdot W^{[1:k]} \cdot h_j + V_R \begin{bmatrix} h_i \\ h_j \end{bmatrix} + b) \quad (17)$$

where σ is a non-linear activation function, $W^{[1:k]} \in \mathbb{R}^{d \times d \times k}$ is a tensor vector with k slices, $V \in \mathbb{R}^{k \times 2d}$ is the weight matrix of a standard neural network and b is the bias. The bilinear tensor product $h_i^T \cdot W^{[1:k]} \cdot h_j$ computes a representation vector $S \in \mathbb{R}^k$ where each slice of the tensor learns a distinct pattern of similarity between the input embeddings.

Figure 2: Overall architecture and workflow of *funcGNN*

3.1.4 Similarity Score Estimation. The final task is to generate the similarity score $\hat{y}_{i,j}$, by passing $NTN_{(h_i, h_j)}$ through multiple layers of fully connected neural networks. A fully connected layer is a non-linear neural network which maps the input of dimension m to a desired output dimension n using multilayered weighted neuron multiplications. From the final layer of the fully connected neural network, one score is calculated which represents the $\hat{y}_{i,j}$. In the training phase, we compare the generated $\hat{y}_{i,j}$ with the actual ground truth $y_{i,j}$ and minimize the *mean square error* (MSE) loss using a gradient-descent strategy. This can be represented as,

$$loss_{mse} = \frac{1}{D} \sum_{(i,j) \in D} (\hat{y}_{i,j} - y_{i,j})^2 \quad (18)$$

where D is the number of graph pairs in the dataset.

3.2 Bottom-Up Approach - Atomic Level Node Comparison

When generating an overall embedding of a graph using the top-down approach, there is the possibility of losing some of the local node information. To overcome this, we also use a *bottom-up approach*. Here, instead of comparing the entire graph embedding we match the similarity between nodes among the graph pair. The idea is to extract atomic node level similarity in an analogous way to the methodology of random walk kernels[36] on graphs. However, graph kernels are computationally expensive and at least of the order $O(n^3)$ [48]. To obtain the node similarities we instead take the inner product of all the pairwise node combinations of the two graphs using their embedding as discussed in Section 3.1.1. The inner product score is later transformed by passing it through a non-linear activation function. To make sure that the graphs are of the same size, we pad the shortest graph in the pair with nodes having zero embedding, i.e. an embedding with zero vector initialization, until both graphs are of the same size. The result of the inner product is thus a similarity matrix, $S \in \mathbb{R}^{N \times N}$, where $N = \max(N_{G_i}, N_{G_j})$ and N_{G_i} denotes the number of nodes in graph G_i .

To efficiently utilise the atomic level node pair similarity we transform the similarity matrix S into a histogram feature vector $H(S) \in \mathbb{R}^b$, where b defines the number of bins. Histograms provide the probability distribution of the node similarities in S and are invariant to the node ordering. This is the same issue which causes the graph isomorphism problem to have high computational complexity. The histogram feature vector $H(S)$ thus obtained is passed to the fully connected layers after normalization and concatenated with $NTN_{(h_i, h_j)}$, to calculate the final similarity score.

4 EXPERIMENTAL DESIGN AND RESULTS OF TRAINING AND EVALUATION

In this section, we describe the training and evaluation of *funcGNN* on a labeled CFG dataset derived from open source Java programs. We describe our experimental setup and present the results obtained from this empirical study.

4.1 The CFG Dataset

For the task of learning program similarity we collected a set of 45 open source Java functions (such as Bubble Sort). These are rich in mathematical operations. Since these programs were all distinct algorithms, their GED values were very high. To overcome this issue, we generated mutants of these programs as a data augmentation method first presented in [34]. A mutant is a structurally modified version of a program, usually created by injecting a fault into it [15], e.g. by changing an operator or relation. By restricting change to one operator, each mutant has a low GED from its original program. Thus we created a dataset having program pairs with both high and low GED values. We generated 4 mutants of each program thus extending our dataset to 225 java functions. To extract the CFGs from these programs we used the Soot framework as shown in Figure 1. We constructed a graph pair dataset of 50625 program pairs in a JSON format which included the node labels, node attributes, edge list and their approximate edit distance using the QAP method. Figure 3 shows the approximate GED values between array division program and two of its mutants. Figure 4 shows the graph size distribution of our dataset. We split our dataset using an

Method	MSE in 10^{-3}
QAP [5]	0.0
VJ [12]	14.41
Hungarian [42]	15.97
HED [14]	8.67
GCN [27]	4.58
GraphSAGE [17]	3.61
funcGNN	1.94

Table 1: Comparison results of the mse error rate

80:20 ratio into training and testing datasets. Figure 5 depicts the approx GED distribution in both the training and testing sets. In our study we chose the approximate GED value as the ground truth value for each graph pair. We transformed these ground truth GED values to the ground truth similarity score S_{GT} by first normalizing them and then passing through an exponential function to map into the interval $(0, 1]$. This can be represented as,

$$norm = 2 \cdot \frac{GED(G_i, G_j)}{|G_i| + |G_j|} \quad (19)$$

$$S_{GT} = e^{-norm} \quad (20)$$

4.2 Results

All experiments were conducted on a single mini workstation: 2.60GHz Intel(R) Xeon(R) CPU E5-2697 v3, 56 CPU cores and 250 GB RAM. To evaluate and compare *funcGNN* with state-of-the-art methods, we used the following two metrics:

- *MSE*: We used the mean square error (MSE) to compute the loss between the predicted score and the ground truth, as defined in equation 18. MSE satisfies the mathematical properties of convexity, symmetry, and differentiability, and is sensitive towards outliers in a dataset. Figure 6 shows the *MSE* loss function curve obtained by *funcGNN* for the training and testing datasets. From figure 6 we can infer that the proposed *funcGNN* model has converged deftly for both the train and test datasets. The performance and convergence behavior of the model indicate that MSE is a good approach for optimizing the *funcGNN* model for learning the similarities between programs. We compared the MSE error value obtained by *funcGNN* with other traditional graph edit distance algorithms and graph neural networks in table 1. Our experiments show that *funcGNN* surpasses the traditional methods and provides a better generalised model for finding similarity between programs.
- *Time*: By time we mean the total time taken by each method to estimate the similarity score of the graph pairs for the entire dataset. Table 2 shows the *time* comparison of *funcGNN* along with other metrics to predict the GED values of all the graph pairs. Since GED estimation is time consuming, we harnessed the power of parallel computing to speed up

Time taken in seconds		
Method	Time	#Parallel Process
QAP [5]	9044.72	1
QAP [5]	405.86	45
VJ [12]	2513.73	45
Hungarian [42]	2546.54	45
HED [14]	9880.18	45
GCN [27]	378.96	1
GraphSAGE [17]	379.24	1
funcGNN	379.81	1

Table 2: GED prediction runtime comparison

the computation. The approximate GEDs for traditional approaches were calculated asynchronously via *ProcessPoolExecutor*⁶ [20] using a pool of 45 concurrent processes. Our results show that *funcGNN* even on serial execution provides faster results than the parallel execution (45 processes) of all the approximation methods.

4.3 Case Studies

We demonstrate three case studies of the predictions made by the proposed *funcGNN* model. All the case study examples are taken from the test dataset and hence unseen by the trained model. In the case study examples, we explain the performance of *funcGNN* when applied on program pairs having high and low similarities. We also provide an example where *funcGNN* has a high error value leading to a wrong prediction. The results of all the case studies are consolidated and presented in Table 3

4.3.1 Case Study 1 : Program pairs with high similarity. We wanted to analyse the ability of *funcGNN* to learn program pairs with high similarity score. The case were two programs will have the highest similarity value will be when they are identical. Hence we randomly chose a program pair (*elementwiseMax_DC_EQ_m3*, *elementwiseMax_DC_EQ_m3*) from our test dataset in which both of the programs are the same. *elementwiseMax_DC_EQ_m3* is an equivalent mutant version of *elementwiseMax* program generated using the methodology described in [34] and outputs the elementwise maximum of two arrays. Figure 7 depicts the control flow graph of *elementwiseMax_DC_EQ_m3* function. Since both the programs in the pair are the same, the ground truth similarity score between them is 1. The proposed *funcGNN* model predicted the similarity value for this pair as 0.9723 with an error of 0.026

4.3.2 Case Study 2 : Program pairs with low similarity. We chose a program pair in the test dataset which had low similarity score. We chose the pair (*heapSort_L_EQ_m4*, *bitwiseOr_L_EQ_m4*) which had a ground truth similarity score of 0.0108. Figure 7 shows the control flow graphs of *heapSort_L_EQ_m4* and *bitwiseOr_L_EQ_m4* functions respectively. The proposed *funcGNN* model predicted the similarity value for this pair as 0.0035 with an error of 0.0073

⁶<https://docs.python.org/3/library/concurrent.futures.html>

Figure 3: Approximate GED values between array division program and two of its mutants

Figure 4: Distribution of graph sizes in the dataset

4.3.3 Case Study 3 : Program pair which received high error rate. Here we demonstrate a program pair in the test dataset which received highest error score between their ground truth and prediction score. The pair (*calVariance*, *countZeros_DC_EQ_m3*) which has a low ground truth similarity of 0.1842. Figure 8 shows the

control flow graphs of *calVariance* and *countZeros_DC_EQ_m3* functions respectively. The proposed funcGNN model predicted the similarity value for this pair as 0.1036 with an error of 0.0806

5 LIMITATIONS AND THREATS TO VALIDITY

There exist multiple factors that could be considered threats to the validity of our results. These include:

- **CFG creation** In this study we have used the open source tool Soot for generating the CFGs. Though the Jimple address format provided by Soot gives a detailed atomic level representation of a program, it can only analyse Java programs. Thus the scope of this project is limited to Java programs and repositories.
- **Data Variability** We had initially taken a small set of Java programs and mutated them to create four sets of similar variants. The reason for this approach was to have examples of programs with small GED values in the dataset. This comes with the drawback that it reduces the variability in the structure of programs in the dataset. However, the idea of this study was to understand the approximate GED among program graph pairs and not its logic or working.
- **Program Size** Since the GED problem is NP-hard, we used individual Java unit functions as our dataset and not the entire Java class file, in order to reduce the number of nodes in

Figure 5: Distribution of Approx GED values in the training and testing sets

Figure 6: Loss function for both training and test sets

Case Study	Program 1	Program 2	Ground Truth	Prediction	Error
1	elementwiseMax_DC_EQ_m3	elementwiseMax_DC_EQ_m3	1.0	0.9732	0.0268
2	heapSort_L_EQ_m4	bitwiseOr_L_EQ_m4	0.0108	0.0035	0.0073
3	calVariance	countZeros_DC_EQ_m3	0.1842	0.1036	0.0806

Table 3: Case study observations from the test dataset

each graph. It would be interesting to see how our approach generalises to predicting the approximate GED of large Java class files.

- **Loss by Approximation** We have employed an approximate value of the actual GED by converting it to a Quadratic Assignment Problem (QAP)[5]. Though this provides an estimate close to the actual GED score, there is always a trade-off when we prioritise computation time over accuracy. Thus

our approach possess all the limitations of the GED calculation used in the QAP approximation method.

- **Backpropagation of Histogram** In our *Bottom-up* approach we extracted the histogram feature representation by performing atomic level node comparison. However, histograms cannot be trained using the backpropagation methodology as there is not a continuous differential function. We use the histogram features just to enhance the global graph features as in [2].

Figure 7: Control flow graphs of `elementwiseMax_DC_EQ_m3`, `heapSort_L_EQ_m4` and `bitwiseOr_L_EQ_m4` functions

6 RELATED WORK

The task of program or code similarity is of fundamental interest in software engineering, and can be traced back to [3]. Program similarity is applied in many software engineering problems, including code plagiarism[11, 54], authorship identification [10, 21], code search [24, 39], clone identification and refactoring [29, 56], and detecting malware patterns [8, 23]. [50] provides a detailed review of numerous methodologies to estimate program similarity and compare code.

There has been some effort to solve the problem of program similarity using traditional machine learning techniques. [19, 26, 32, 41]. However all these studies used features hand-crafted manually by domain experts, which is expensive in terms of time and expertise. The use of deep learning techniques which avoid this feature engineering approach, was applied for solving program similarity in [33]. One drawback of deep learning is the amount of tagged data it requires for training, which might be difficult to obtain in the field of software engineering. [34] demonstrates the use of equivalent mutants as a data augmentation method on source code, to alleviate the data crunch problem. In our approach we have used this data augmentation methodology for creating program pairs with low GED score.

Another limitation of deep learning models is that, they are mainly trained on the shallow textual structure of a program (syntax), and can miss out on semantic features [1]. The study in [1] suggests that using graphs we can represent both the syntactic and semantic structure of code, and it demonstrates the effectiveness of graph neural networks for better program analysis.

Graphs, especially control flow graphs, have been used extensively to solve many software engineering problems [13, 22, 28, 35, 40, 49]. In [55], labeled CFGs were analysed using deep learning techniques for learning code semantic similarity. However, the input to the deep learning model in this study was a hand-crafted feature matrix thus restricting the models's capability of inferring the semantics of the graph. Graph neural networks (GNNs) have emerged as a successful class of neural network models capable of learning of graph representations effectively [16, 17, 27, 45, 51]. In our approach we have harnessed the graph learning capability of GNNs to analyse labeled CFG graphs for program similarity.

The most similar work to our approach is [2] where the authors propose SimGNN for graph similarity. The main differences between these two studies are: (i) the neighbourhood aggregation method (GCN[27] in SimGNN and GraphSAGE[17] in ours), (ii) the choice of hyperparameters and the activation functions used, and (iii) the dataset used. Regarding (iii), in [2] the authors evaluate their model on unlabeled program dependency graphs (PDG) of C

Figure 8: Control flow graphs of calVariance and countZeros_DC_EQ_m3 functions

programs. All the unlabeled nodes in their approach were initialised with the same label, leading to each node having the same initial embedding representation. Hence the only code features learned were the data flow features and not the node features (as in our approach). In our approach we trained on labeled CFGs where each node in the CFG was labeled with an atomic program statement. Each such atomic operation was initialised with an unique embedding, thus providing the learned model with richer program structure. Also the dataset of [2] consisted of small programs each restricted to a maximum value of 10. In our approach, the model was trained on larger program graphs with an average size of 19.76 and a maximum size of 72 (see Figure 4).

Other interesting contributions similar to ours are [31, 52], where labeled CFGs are analysed to finding binary code similarities and to check the presence of any known security vulnerabilities in it. However, here each node in the CFG dataset represents multiple attributes (such as mov, lea, cmp and jbe all in one node). This makes it difficult to calculate an appropriate embedding to that node which summarizes all of its operations. In our approach, each atomic operation in the program is assigned a new node to achieve a better node embedding.

7 CONCLUSION

In this paper, we have studied the problem of program similarity using graph edit distance and proposed *funcGNN* a graph neural network approach to estimating the GED. To characterise the semantics and logic of the program we analyse its control flow representation. *funcGNN* inherits the inductive and node-order invariant properties of graph convolutional networks and uses it to create a semantically rich embedding for each node in a CFG. The evaluation study carried out in our research shows that *funcGNN* is capable of estimating the approximate graph edit distance of unseen program pairs with very low error rate and is computationally efficient. We have discussed the limitations and drawbacks of our approach, and we see potential improvements in future. We will also consider how to apply *funcGNN* for solving related software engineering challenges such as clone refactoring, for which finding program similarity is crucial.

8 ACKNOWLEDGMENTS

The authors gratefully acknowledge financial support from the ITEA3 TESTOMAT Project 16032 and Ericsson AB.

REFERENCES

- [1] Miltiadis Allamanis, Marc Brockschmidt, and Mahmoud Khademi. 2017. Learning to represent programs with graphs. *arXiv preprint arXiv:1711.00740* (2017).
- [2] Yunsheng Bai, Hao Ding, Song Bian, Ting Chen, Yizhou Sun, and Wei Wang. 2019. Simgnn: A neural network approach to fast graph similarity computation. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 384–392.
- [3] Hal I. Berghel and David L. Sallach. 1984. Measurements of program similarity in identical task environments. *ACM SIGPLAN Notices* 19, 8 (1984), 65–76.
- [4] David B. Blumenthal and Johann Gamper. 2018. On the exact computation of the graph edit distance. *Pattern Recognition Letters* (2018).
- [5] Sebastien Bouteux, Luc Brun, Vincenzo Carletti, Pasquale Foggia, Benoit Gaüzère, and Mario Vento. 2017. Graph edit distance as a quadratic assignment problem. *Pattern Recognition Letters* 87 (2017), 38–46.
- [6] Horst Bunke. 1997. On a relation between graph edit distance and maximum common subgraph. *Pattern Recognition Letters* 18, 8 (1997), 689–694.
- [7] Horst Bunke and Kim Shearer. 1998. A graph distance metric based on the maximal common subgraph. *Pattern recognition letters* 19, 3–4 (1998), 255–259.
- [8] Silvio Cesare and Yang Xiang. 2011. Malware variant detection using similarity search over sets of control flow graphs. In *2011 IEEE 10th International Conference on Trust, Security and Privacy in Computing and Communications*. IEEE, 181–189.
- [9] William W. Cohen, Pradeep Ravikumar, Stephen E. Fienberg, et al. 2003. A Comparison of String Distance Metrics for Name-Matching Tasks.. In *IIWeb*, Vol. 2003. 73–78.
- [10] Chandan Dasgupta. 2010. That is not my program: Investigating the relation between program comprehension and program authorship. In *Proceedings of the 48th Annual Southeast Regional Conference*. 1–4.
- [11] Jinan AW. Faidhi and Stuart K. Robinson. 1987. An empirical approach for detecting program similarity and plagiarism within a university programming environment. *Computers & Education* 11, 1 (1987), 11–19.
- [12] Stefan Fankhauser, Kaspar Riesen, and Horst Bunke. 2011. Speeding up graph edit distance computation through fast bipartite matching. In *International Workshop on Graph-Based Representations in Pattern Recognition*. Springer, 102–111.
- [13] Qian Feng, Rundong Zhou, Chengcheng Xu, Yao Cheng, Brian Testa, and Heng Yin. 2016. Scalable graph-based bug search for firmware images. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 480–491.
- [14] Andreas Fischer, Kaspar Riesen, and Horst Bunke. 2017. Improved quadratic time approximation of graph edit distance by combining Hausdorff matching and greedy assignment. *Pattern Recognition Letters* 87 (2017), 55–62.
- [15] Robert Geist, A. Jefferson Offutt, and Frederick C. Harris Jr. 1992. Estimation and enhancement of real-time software reliability through mutation analysis. *IEEE Trans. Comput.* 5 (1992), 550–558.
- [16] Marco Gori, Gabriele Monfardini, and Franco Scarselli. 2005. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, Vol. 2. IEEE, 729–734.
- [17] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*.

- 1024–1034.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. In *European conference on computer vision*. Springer, 630–645.
- [19] Kenneth Hoste, Aashish Phansalkar, Lieven Eeckhout, Andy Georges, Lizy K John, and Koen De Bosschere. 2006. Performance prediction based on inherent program similarity. In *2006 International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 114–122.
- [20] John Hunt. 2019. Futures. In *Advanced Guide to Python 3 Programming*. Springer, 395–405.
- [21] Vaibhavi Kalgutkar, Ratinder Kaur, Hugo Gonzalez, Natalia Stakhanova, and Alina Matyukhina. 2019. Code authorship attribution: Methods and challenges. *ACM Computing Surveys (CSUR)* 52, 1 (2019), 1–36.
- [22] Upulee Kanewala, James M Bieman, and Asa Ben-Hur. 2016. Predicting metamorphic relations for testing scientific software: a machine learning approach using graph kernels. *Software testing, verification and reliability* 26, 3 (2016), 245–269.
- [23] Abhishek Karnik, Suchandra Goswami, and Ratan Guha. 2007. Detecting obfuscated viruses using cosine similarity analysis. In *First Asia International Conference on Modelling & Simulation (AMS'07)*. IEEE, 165–170.
- [24] Iman Keivanloo, Juergen Rilling, and Ying Zou. 2014. Spotting working code examples. In *Proceedings of the 36th International Conference on Software Engineering*. 664–675.
- [25] Mohamed A Khamsi and William A Kirk. 2011. *An introduction to metric spaces and fixed point theory*. Vol. 53. John Wiley & Sons.
- [26] Yesol Kim, Jonghyuk Park, Seong-je Cho, Yunmook Nah, Sangchul Han, and Minkyu Park. 2015. Machine learning-based software classification scheme for efficient program similarity analysis. In *Proceedings of the 2015 Conference on research in adaptive and convergent systems*. 114–118.
- [27] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [28] Jens Krinke. 2006. Mining control flow graphs for crosscutting concerns. In *2006 13th Working Conference on Reverse Engineering*. IEEE, 334–342.
- [29] Giri Panamoottil Krishnan and Nikolaos Tsantalis. 2014. Unification and refactoring of clones. In *2014 Software Evolution Week-IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering (CSMR-WCRE)*. IEEE, 104–113.
- [30] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [31] Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. 2019. Graph matching networks for learning the similarity of graph structured objects. *arXiv preprint arXiv:1904.12787* (2019).
- [32] Jonathan I Maletic and Andrian Marcus. 2000. Using latent semantic analysis to identify similarities in source code to support program understanding. In *Proceedings 12th IEEE International Conference on Tools with Artificial Intelligence. ICTAI 2000*. IEEE, 46–53.
- [33] Niccolò Marastoni, Roberto Giacobazzi, and Mila Dalla Preda. 2018. A deep learning approach to program similarity. In *Proceedings of the 1st International Workshop on Machine Learning and Software Engineering in Symbiosis*. 26–35.
- [34] Aravind Nair, Karl Meinke, and Sigrid Eldh. 2019. Leveraging mutants for automatic prediction of metamorphic relations using machine learning. In *Proceedings of the 3rd ACM SIGSOFT International Workshop on Machine Learning Techniques for Software Quality Evaluation*. 1–6.
- [35] Animesh Nandi, Atri Mandal, Shubham Atreja, Gargi B Dasgupta, and Subhrajit Bhattacharya. 2016. Anomaly detection using program control flow graph mining from execution logs. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 215–224.
- [36] Michel Neuhaus and Horst Bunke. 2006. A random walk kernel derived from graph edit distance. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 191–199.
- [37] Michel Neuhaus and Horst Bunke. 2007. A quadratic programming approach to the graph edit distance problem. In *International Workshop on Graph-Based Representations in Pattern Recognition*. Springer, 92–102.
- [38] Michel Neuhaus, Kaspar Riesen, and Horst Bunke. 2006. Fast suboptimal algorithms for the computation of graph edit distance. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 163–172.
- [39] Haoran Niu, Iman Keivanloo, and Ying Zou. 2017. Learning to rank code examples for code search engines. *Empirical Software Engineering* 22, 1 (2017), 259–291.
- [40] Anh Viet Phan, Minh Le Nguyen, and Lam Thu Bui. 2017. Convolutional neural networks over control flow graphs for software defect prediction. In *2017 IEEE 29th International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 45–52.
- [41] Aashish Phansalkar, Ajay Joshi, Lieven Eeckhout, and Lizy Kurian John. 2005. Measuring program similarity: Experiments with SPEC CPU benchmark suites. In *IEEE International Symposium on Performance Analysis of Systems and Software, 2005. ISPASS 2005*. IEEE, 10–20.
- [42] Kaspar Riesen and Horst Bunke. 2009. Approximate graph edit distance computation by means of bipartite graph matching. *Image and Vision computing* 27, 7 (2009), 950–959.
- [43] Kaspar Riesen, Andreas Fischer, and Horst Bunke. 2014. Improving approximate graph edit distance using genetic algorithms. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 63–72.
- [44] Eric Sven Ristad and Peter N Yianilos. 1998. Learning string-edit distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 5 (1998), 522–532.
- [45] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE Transactions on Neural Networks* 20, 1 (2008), 61–80.
- [46] Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In *Advances in neural information processing systems*. 926–934.
- [47] Raja Vallee-Rai and Laurie J Hendren. 1998. Jimple: Simplifying Java bytecode for analyses and transformations. (1998).
- [48] SVN Vishwanathan, Karsten M Borgwardt, Nicol N Schraudolph, et al. 2006. Fast computation of graph kernels. In *NIPS*, Vol. 19. 131–138.
- [49] Milena Vujosević-Janičić, Mladen Nikolić, Dušan Tošić, and Viktor Kuncak. 2013. Software verification and graph similarity for automated evaluation of students' assignments. *Information and Software Technology* 55, 6 (2013), 1004–1016.
- [50] Andrew Walenstein, Mohammad El-Ramly, James R Cordy, William S Evans, Kiarash Mahdavi, Markus Pizka, Ganesan Ramalingam, and Jürgen Wolff von Gudenberg. 2007. Similarity in programs. In *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [51] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018).
- [52] Xiaojun Xu, Chang Liu, Qian Feng, Heng Yin, Le Song, and Dawn Song. 2017. Neural network-based graph embedding for cross-platform binary code similarity detection. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 363–376.
- [53] Zhiping Zeng, Anthony KH Tung, Jianyong Wang, Jianhua Feng, and Lizhu Zhou. 2009. Comparing stars: On approximating graph edit distance. *Proceedings of the VLDB Endowment* 2, 1 (2009), 25–36.
- [54] Fangfang Zhang, Dinghao Wu, Peng Liu, and Sencun Zhu. 2014. Program logic based software plagiarism detection. In *2014 IEEE 25th International Symposium on Software Reliability Engineering*. IEEE, 66–77.
- [55] Gang Zhao and Jeff Huang. 2018. Deepsim: deep learning code functional similarity. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 141–151.
- [56] Minhaz Fahim Zibran and Chanchal Kumar Roy. 2013. Conflict-aware optimal scheduling of prioritised code clone refactoring. *IET software* 7, 3 (2013), 167–186.