

Dynamic Neural Control Flow Execution: an Agent-Based Deep Equilibrium Approach for Binary Vulnerability Detection

Litao Li
litao.li@queensu.ca
Queen's University
Kingston, Ontario, Canada

Steven H. H. Ding
steven.h.ding@mcgill.ca
McGill University
Montreal, Quebec, Canada

Andrew Walenstein
walenste@ieee.org
BlackBerry Limited
Waterloo, Ontario, Canada

Philippe Charland
philippe.charland@drdc-rddc.gc.ca
Mission Critical Cyber Security Section
Defence R&D Canada
Quebec City, Quebec, Canada

Benjamin C. M. Fung
ben.fung@mcgill.ca
McGill University
Montreal, Quebec, Canada

ABSTRACT

Software vulnerabilities are a challenge in cybersecurity. Manual security patches are often difficult and slow to be deployed, while new vulnerabilities are created. Binary code vulnerability detection is less studied and more complex compared to source code, and this has important practical implications. Deep learning has become an efficient and powerful tool in the security domain, where it provides end-to-end and accurate prediction. Modern deep learning approaches learn the program semantics through sequence and graph neural networks, using various intermediate representation of programs, such as abstract syntax trees (AST) or control flow graphs (CFG). Due to the complex nature of program execution, the output of an execution depends on the many program states and inputs. Also, a CFG generated from static analysis can be an overestimation of the true program flow. Moreover, the size of programs often does not allow a graph neural network with fixed layers to aggregate global information. To address these issues, we propose DeepEXE, an agent-based implicit neural network that mimics the execution path of a program. We use reinforcement learning to enhance the branching decision at every program state transition and create a dynamic environment to learn the dependency between a vulnerability and certain program states. An implicitly defined neural network enables nearly infinite state transitions until convergence, which captures the structural information at a higher level. The experiments are conducted on two semi-synthetic and two real-world datasets. We show that DeepEXE is an accurate and efficient method and outperforms the state-of-the-art vulnerability detection methods.

CCS CONCEPTS

• **Security and privacy** → **Software reverse engineering**; • **Computing methodologies** → **Neural networks**.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the national government of Canada. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Request permissions from owner/author(s).

CIKM '24, October 21–25, 2024, Boise, ID, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0436-9/24/10

<https://doi.org/10.1145/3627673.3679726>

KEYWORDS

Neural Networks, Cybersecurity, Vulnerability Detection, Implicit Deep Learning, Graph Neural Network, Reinforcement Learning, Binary Code, Control Flow Graph

ACM Reference Format:

Litao Li, Steven H. H. Ding, Andrew Walenstein, Philippe Charland, and Benjamin C. M. Fung. 2024. Dynamic Neural Control Flow Execution: an Agent-Based Deep Equilibrium Approach for Binary Vulnerability Detection. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3627673.3679726>

1 INTRODUCTION

Software vulnerabilities have been an ongoing challenge in the cybersecurity domain. It is an inevitable problem, as the scale of software grows in complexity. Many malicious cyber attacks exploit vulnerabilities within systems and can cause tremendous economic and security damages. Often, the security analysts cannot even patch vulnerabilities fast enough, as new ones are created [2, 18]. Common Vulnerability Exposures (CVE) show that the total number of vulnerabilities more than doubled from 2016 to 2017 and it continued to increase throughout the recent years¹. Many traditional static and dynamic analysis methods are manually expensive and inefficient. This motivates automated and end-to-end approaches, such as neural networks.

Vulnerabilities can be detected at either the source code or binary code level. Source code provides much more meaningful semantics, syntax, and structures, which in turn help both analysts and machine learning models to track vulnerabilities. Existing methods at the source code level are accurate and capable of finding complex vulnerabilities [24, 37]. For binary code, as much information is lost during the compilation process, it is much harder to detect vulnerabilities. Moreover, the absence of the original source is a practical problem under many circumstances, such as third-party or off-the-shelf programs. Binary code is best analyzed as assembly code, a form of intermediate representation that provides analysts readable content. Assembly code contains instructions that provide some semantics and structures of the program. In this paper, we

¹Statistics on Common Vulnerabilities and Exposures (CVE) Details

are only interested in binary code vulnerability detection, as it is still a prevalent challenge in the security domain.

Deep learning methods aim to learn the latent representation of a piece of binary code for classification. Existing works for binary code learning can be categorized into two main streams. The first approach focuses on text-based representation learning to extract the token semantics. The instructions are broken down and embedded into vectors through some unsupervised learning such as Word2Vec [39], then these vectors are fed into a sequential deep learning model for classification. **Instruction2Vec** [33], **HAN-BSVD** [53], and **BVDetector** [46] all use this semantic-based approach for detection. The second method involves collecting and aggregating structural information at a higher level. Usually, CFGs are parsed from the assembly code basic blocks, which create dependencies between different blocks of code. This is crucial in vulnerability detection, since programs are complex and hierarchical, and vulnerabilities are often triggered in specific program states. Using only the semantics of instruction tokens are often insufficient. **Gemini** [52], **Diff** [38], **Order** [54], **InnerEye** [57], and **BinDeep** [45] all use graph-based methods for binary code structure embedding.

Unfortunately, there are major drawbacks to either approach that can hinder the performance or scalability of the model. The more obvious disadvantage is the scalability when large programs are present. Semantic-based approaches usually introduce a maximum input length, in order to prevent vanishing gradient, especially for large and deep sequence models. Structure-based approaches perform graph neural network (GNN) for aggregating node information. The number of layers dictates the receptive field of the model by performing k -hop message passing, thus limiting the amount of global information that can be learned. Both of them need to carefully manage the memory footprint during training. The other drawback is the absence of modelling how programs naturally run. Unlike natural language, programs are executed dynamically. The state of a program can be different, depending on the input and its previous states. By using fixed graph learning techniques, the dynamic nature of the program structure is difficult to capture and thus lead to undesired performance.

Given assembly code, one has to respectively find a program execution path that can potentially yield the same final program state. In general, a sound and complete static analysis method generates a representation of the code (i.e., CFG) with overestimation. This means paths created in a graph can potentially never execute. Ideally, symbolic execution [6, 27] is one of the formal methods that enable one to compare and verify all the possible paths through equivalence checking. However, its applicability is limited, as it requires storing all the possible program states associated with all possible execution paths. This will cause the path explosion problem [51], especially for large functions with loops. Existing works try to address the pathfinding problem statically from an incomplete view, focusing on partial or local structures. For example, **DeepBinDiff** [16] and **InnerEye** [57] match the CFGs based on semi-exhaustive path comparison, which is not scalable, and also misses the iterative graph learning. **Genius** [52], **BinGo** [9], and **Tracelet** [13] use partial path matching, which lacks robustness when programs are easily altered through artificial means. **BinaryAI** [54] uses graph convolution for message passing. However,

this approach does not consider mutually exclusive dependencies among edges, covering invalid paths. The message passing mechanism also assumes a static adjacency matrix, which lacks high-level guidance from a global state. The current research in this domain lacks a dedicated way to simulate the program state transitions along the guided valid execution path, with a focus on a higher order of node neighbourhood proximity.

Inspired by symbolic execution for path-finding, we propose a neural network model, DeepEXE, which mimics a program state-guided execution process over the CFG to detect binary code vulnerabilities at the function or file level. DeepEXE relies on an execution agent that simulates and learns which direction to take, resulting in simulated paths across different epochs. The combined node embedding represents the program state, and the branching actions guiding the program flow are based on the program state and code semantics of the current node. DeepEXE leverages the implicit neural network paradigm, where only the final program state is stored before back-propagation. This enables a large simulation step over the execution flow. Compared to the existing methods with only local or partial graph information, DeepEXE enables modelling on the highest global-level view over the execution path. Our contributions are as follows:

- We propose DeepEXE, a neural program execution model over a CFG for binary vulnerability detection. It simulates a semantic-guided decision process for stepping through a given function's CFG.
- To simulate the program execution steps over the graph, we propose a learning agent for making branching decisions with an implicit neural network structure for program state transitions. It enables modelling program semantics on a higher level views over the execution path.
- To address the scalability and limited receptive field of graph neural networks, we use the implicit deep learning paradigm for nearly infinite message passing, which significantly enables global information aggregation in the graph and reduces the memory footprint.
- We conduct experiments on two semi-synthetic datasets and two real world vulnerability datasets. We compare our methods against several state-of-the-art approaches and show that DeepEXE can consistently outperform the baselines in all scenarios.

2 RELATED WORK

Vulnerability Detection While vulnerability detection can be conducted at either the source code or binary code level, we will discuss them together, since most methods can be applied to both levels, with some modifications. Machine learning-based (non-deep learning) methods involve the manual extraction of metrics and the input of these metrics as features [22, 44]. The metrics can be multi-level and leverage the complexity characteristics of a program, such as the number of nested loops within a function. Manual feature extraction is more expensive and requires expert knowledge. Also, the features need to be constantly updated to accommodate changes in the codebase. Text-based deep learning is very popular for source code vulnerability detection, where different granularity

levels can be leveraged in order to obtain text features or embeddings. Li et al. group tokens based on semantics and syntax into slices or gadgets [36, 37, 56], and feed them into a LSTM model. For binary code, Instruction2Vec [34] and Bin2img [33] use instruction embedding as a preprocessing step. Similar to Word2Vec, the embedding contains contextual dependency and can be used to detect vulnerabilities at a later stage, which is a 1D CNN model. These models solely focus on the semantics of the tokens, where the structural information is omitted. There are several GNN models at the source code that use different graphs that can be parsed from source code, such as abstract syntax trees, data dependence graphs, and control flow graphs [8, 41, 55]. For GNN message passing, there are multiple styles that we will discuss next.

Graph Neural Networks and Implicit Models In binary code, GNN methods aim at learning the structures by first parsing the assembly code into control flow graphs and performing message passing. There are multiple variants related to graph neural networks. The pioneer works of graph neural networks are mostly associated with recurrent graph neural networks [11, 19, 20, 35, 43], where the node representations are aggregated with a fixed set of parameters. Convolutional graph neural networks [23, 29, 47] expand the GNN by using multiple layers with different parameters. This approach addresses the cyclic mutual dependencies architecturally [50] and is more efficient and powerful. However, GNNs struggle to capture long-range dependencies in large graphs, due to the finite number of message passing iterations. One potential solution is the recently studied implicit neural networks. The implicit learning paradigm is different from traditional deep learning, as it solves the solution for a given equilibrium problem, which is formulated as an nearly infinite layer network. Implicit models have previously shown success in domains such as sequence learning [5], physics engine [14], and graph neural networks [21].

3 PRELIMINARIES

CFGs and Basic Blocks The input of the model is a binary file in assembly code. The assembly functions and their CFGs are both obtained from the IDA Pro disassembler². Each function is regarded as a graph \mathcal{G} that contains segmented code blocks called basic blocks, which are sequences of instructions without any jump or call to other blocks. As the input to the neural network, a graph $\mathcal{G} = (V, A)$ has the blocks $V \in \mathbb{R}^{n \times v}$ with n nodes, v tokens, and the adjacency matrix $A \in \mathbb{R}^{n \times n}$. A defines all directed edges within the graph and is obtained by extracting call statements between the blocks. Note that A has 0 across the diagonal element and is non-symmetrical. Moreover, we apply the re-normalization trick to A [29], in order to prevent numerical instabilities during deep network training. For file level classification, we merge the function graphs as a whole, based on the function call information. Moreover, additional information, such as comments and names, are removed. The basic block V only contains operations and operands of instructions.

Graph Neural Network GNN is a topological learning technique for input data with graph structures. A graph is represented as $\mathcal{G} = (V, E)$ that contains $n := |V|$ nodes and $e := |E|$ edges. An edge

$E_{ij} := (V_i, V_j)$ represents the directed or un-directed connection between node (i, j) . In practice, the edge information is represented in the form of an adjacency matrix $A \in \mathbb{R}^{n \times n}$. Generally, one can obtain some initial node embedding $U \in \mathbb{R}^{n \times h}$ before feeding into the network. The message passing (i.e. node aggregation) is performed at each GNN layer as follows:

$$X^{t+1} = \phi(X^t W^t A^t) \quad (1)$$

where $W^t \in \mathbb{R}^{h \times h}$ is a trainable parameter at layer t . Each message passing step aggregates 1-hop neighbour information into the current node given that an edge exists in A . The final node vector X^T then learns the topological information from all T -hop away neighbours. In case of graph classification, a pooling layer such as add pooling can be used to obtain the graph embedding G :

$$G = \sum_i^n X_{i,j}^T, \forall j = 1, \dots, h \quad (2)$$

REINFORCE Algorithm Reinforcement learning is a class of algorithms that specify the actions within an environment that optimizes the reward r . In particular, the REINFORCE algorithm [49] is a form of policy gradient algorithm that computes the stochastic gradient with respect to the reward. It involves a state s that can be obtained from a neural network, an agent a that specifies the action space \mathcal{A} , and a policy $\pi(a|s)$ that takes the action a given a state s with probabilities. Usually, the policy is randomly initialized and the algorithm iterates through epochs, where backpropagation is performed at each epoch to update the policy in the context of a neural network setup.

Problem Statement We define several neural network modules within our architecture $F = (F_S, F_I, F_A)$, where F_S is the sequential model for semantics embedding, F_I is the implicit graph neural network model for structure and node embedding, and F_A is the reinforcement learning agent for dynamic pathing optimizer, given certain program states. The goal is to predict whether each function contains a vulnerability. Given the input graph $\mathcal{G} = (V, A)$, our model learns several levels of information and aggregates them together for the final output of the model, which is a binary classification score $F : \mathcal{G} \rightarrow \hat{y} \in \mathbb{R}$. Formally, we define the following learning task parameterized by θ :

$$\hat{y} = \operatorname{argmax}_{y' \in \{0,1\}} F_\theta(y' | \mathcal{G}, \theta); \theta = \operatorname{argmax}_\theta F_\theta(y' = y | \mathcal{G}, \theta) \quad (3)$$

4 NEURAL CONTROL FLOW EXECUTION

We design the DeepEXE architecture with semantic-driven and execution-guided principles. We show the overall architecture including the input preprocessing, semantics learning, state transition, and predictions and training in Figure 1.

CFGs extracted from disassembly contain crucial information about the program logic and paths, which dictates the outputs and functionalities of assembly code. An important characteristic to differentiate CFGs from graphs in other domains, such as social networks or chemistry, is that node states should be dependent on the execution logic. Programs are executed following specific orders based on the dependencies among the edges conditioned by the

²IDA Pro

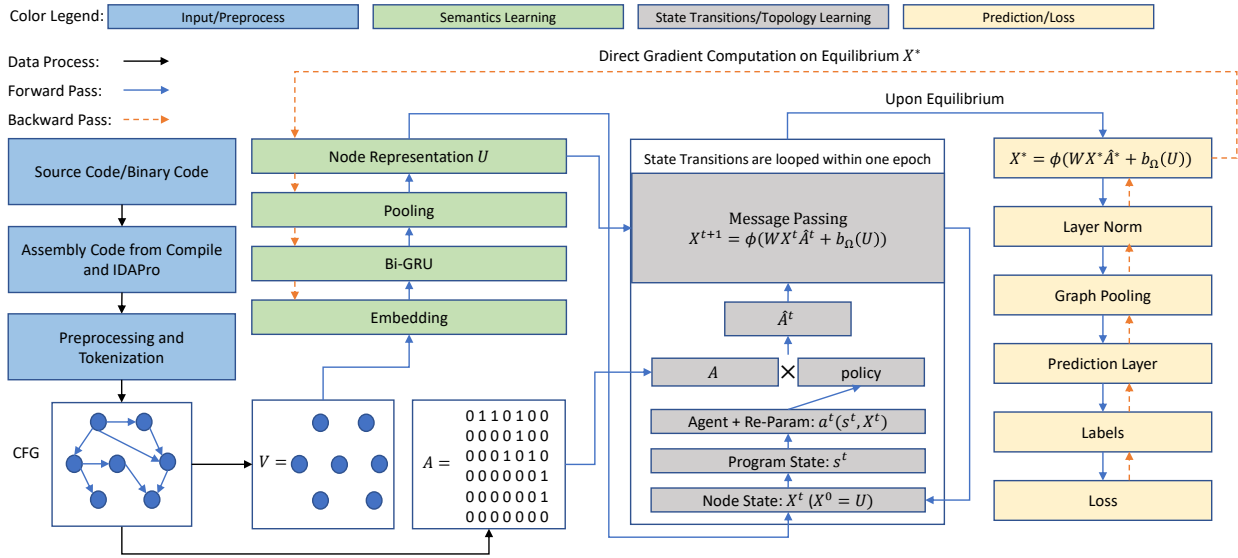


Figure 1: Overall architecture of DeepEXE. Four major segments of our model include input preprocessing, node embedding through sequential semantics, state transition and structure learning, and prediction and training. DeepEXE combines the local instruction semantics with high-level topological information, where program dependencies are captured through the use of a REINFORCE agent and a GNN with much larger receptive field.

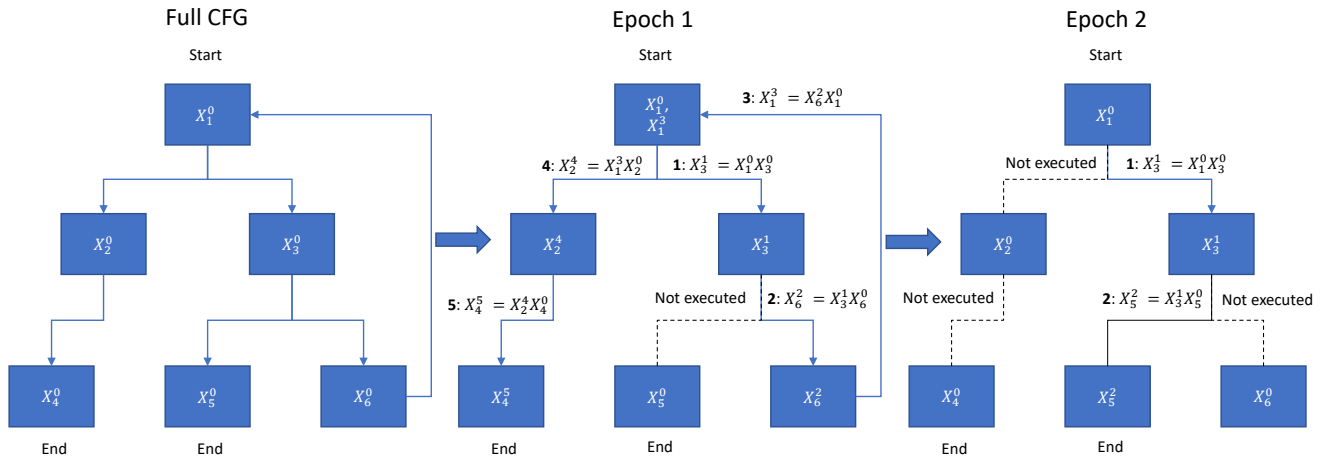


Figure 2: In each epoch, the model simulates one execution session with a specific execution path consisting of multiple steps. At step i , the executor chooses the most likely branch for node j to move next based on program state and node semantics. This is one execution with a loop (Epoch 1) and one without (Epoch 2). The model then updates the program state by combining the next node’s code semantics.

program state, where the results and semantics can substantially differ when orders vary.

We borrow the idea of symbolic execution [6] and create a neural CFG executor. A training epoch contains a full iteration of the executive session, which corresponds to a concrete execution path.

Note that each epoch can have completely different execution paths, as the model learns.

An example of the learning process is shown in Figure 2. The training consists of many training epochs. In Figure 2, the path for Epoch 1 goes into a loop, while Epoch 2 directly goes into the exit point. The execution agent performs multiple steps within an epoch. It starts from the entry node, then transitions to other possible nodes in each step. The decision on which branch to select depends on the program state X , and X_j^t indicates the updated program state at step i for node j . After jumping to the next node, the agent updates the program state and repeats the decision process until it reaches an equilibrium state.

4.1 Token Semantics

In this section, we discuss the preprocessing, embedding, and sequential learning task. A basic block contains a stream of instructions, which can be further broken down into operations and operands, and be tokenized. We treat the entire block as one sentence and apply a subword and unigram [30, 31] model for the token encoding, which mitigates the out-of-vocabulary problem. Assembly code is compiler dependent and can easily result in out-of-vocabulary (OOV) tokens. A way to address the OOV issue is to break down the tokens into characters for encoding. Even with a fixed vocabulary size, unseen tokens can be encoded by matching the subword to their closest known tokens. Moreover, it is not language dependent and can be trained from scratch very efficiently. We increase the subspace representation power by simply applying an embedding layer $E : \mathbf{V} \rightarrow \mathbb{R}^{n \times r \times h}$ next, where r denotes the sequence length and h denotes the hidden dimension. Note that we use h as the hidden dimension throughout the paper for simplicity, but different dimensions can be used for any layers in practice. The sequential model used in this task is a bi-directional GRU [10]. The output of the GRU layer $U \in \mathbb{R}^{n \times r \times h}$ further embeds the token semantics by taking contextual information into account. In order to obtain a representation for the entire basic block, a maximum or average pooling along the time dimension is used to compute $U \in \mathbb{R}^{n \times h}$ for block embedding.

4.2 Program State Guided Execution and Functional Representation

Program State Initial node representation U establishes the semantics within basic blocks, but it is not sufficient to simply globally aggregate U for a high-level representation of the graph. In this regard, a reinforcement agent $a^t(s^{t-1})$ that decides the next execution path is defined, given the previous program state s^{t-1} . Unlike traditional neural networks that perform forward and backward pass one at a time, our approach internally loops through multiple states t within a training epoch. We define the program state as a linear transformation of the node state X^t , where $X^0 = U$, and some trainable parameter $W_s \in \mathbb{R}^{h \times 1}$:

$$s^t = \sigma(X^t W_s) \quad (4)$$

Agent Reparameterization Due to the backpropagation algorithm, categorical variables are hard to train in this stochastic environment in the neural network. This layer effectively becomes non-differentiable when using normal sampling process such as argmax.

A solution is to use the Gumbel softmax [26] to re-parameterize the state while maintaining the ability to backpropagate efficiently during training. Gumbel softmax is a continuous and differentiable distribution that can sample categorical distribution, it is given by:

$$z_i^t = \frac{\exp((\log(s_i^{t-1}) + g_i)/\tau)}{\sum_j^k \exp((\log(s_j^{t-1}) + g_j)/\tau)}, \text{ for } i = 1, \dots, k \quad (5)$$

where z_i^t is the sample drawn from the state, $g_i \sim \text{Gumbel}(0, 1)$ are samples drawn i.i.d from the Gumbel distribution, and τ is the temperature controlling the discreteness of the new samples. Gumbel softmax works better with a lower value for $\tau \in [0, \infty]$ as it approaches to argmax smoothly, whereas setting a large value makes the samples become uniform.

Adjacency Matrix Update In each state update, the agent walks through the graph with updated program state to capture the intermediate execution path that leads to certain results. We have the flexibility to design the agent to be either hard or soft. A soft agent $a^t = z^t$ preserves the probabilities drawn from Gumbel softmax, which implies that a program information can flow in different execution paths at the same time based on the probabilities $\sum_i z_i^t = 1$. A hard agent mimics the execution path and is one-hot, leading to one strictly one execution at a time. The agent $a^t \in \mathbb{R}^{n \times 1}$ is then used to select a path and generate the state-dependent adjacency matrix \tilde{A}^t , which is updated as: $\tilde{A}^t = A a^t$.

4.3 Executor Stepping Via Implicit GNN

Implicit GNN With the updated adjacency matrix from the agent, one can perform graph neural network on the CFG to aggregate neighbour information into the nodes. However, assembly code can be large for various reasons. For example, a GCC compiler can use an optimization level that minimizes the execution size and reduces the size of CFGs. While GNN is a suitable approach to learn the structural dependency of a function, it requires a pre-defined number of layers, where each layer usually performs 1-hop message passing. Intuitively, the vanilla GNNs do not scale well with large graphs and can fail to capture global information. The dependency between further nodes can be crucial to understand the overall semantics of a program. Such long range dependency is difficult to capture with longer edges. To alleviate the above stated problem, we perform the program state transitions in an implicitly defined style. In general, the transition at state t can be written as an implicit form of the GNN layer:

$$X^{t+1} = \phi(X^t W^t A^t + U) \quad (6)$$

$$y^t = f_\psi(X^*) \quad (7)$$

Such form of layer does not explicitly output a vector to be fed into the next layer. Instead, it uses a fixed point iteration in equation 6 that aims to find the equilibrium vector state X^* as $t \rightarrow \infty$. The equilibrium state is then used for the prediction task in equation 7, where f_ψ is an output function parameterized by ψ for the desired classification task. With the reinforcement agent embedded in the updated adjacency matrix $\tilde{A}^* = \tilde{A}^t : t \rightarrow \infty$, our equilibrium solution is formulated as follows:

$$X^* = \phi(X^* W \tilde{A}^* + b_\Omega(U)) \quad (8)$$

$W \in \mathbb{R}^{h \times h}$ and $\Omega \in \mathbb{R}^{h \times h}$ are parameters, and U is the initial node feature. Note that only a single layer is required to produce the updated node representation X iteratively instead of multiple stacking layers. We also inject U into the equation through some affine transformation b_Ω . This ensures that original node semantics is preserved throughout the iterations when solving for the fixed point [5].

Fixed Point Acceleration Although the equilibrium point can be obtained from iterating equation 8 infinitely, it is not the most efficient and stable method for convergence. More importantly, it does not guarantee convergence. Anderson acceleration [48] is an accelerated algorithm for finding fixed points. Given a function f to solve, which is equation 8 in our case, we define (1) $m_k = \min\{m, t\}$ as the parameter for controlling past iteration memory by setting m to any positive integer; (2) $g(x) = f(x) - x$ as the residual with the matrix $G_t = [g_{t-m_t}, \dots, g_t]$. The root solving process using Anderson acceleration is formulated as:

$$\alpha_t = \operatorname{argmin}_\alpha \|G_t \alpha\|_2, \quad (9)$$

$$\text{where } \alpha = (\alpha_0, \dots, \alpha_{m_t}) \in \mathbb{R}^{m_t+1} : \sum_{i=0}^{m_t} \alpha_i = 1$$

$$x^{t+1} = \sum_{i=0}^{m_t} (\alpha_i)_i f(x_{t-m_t+i}) \quad (10)$$

Instead of computing for x^{t+1} directly from x^t , Anderson acceleration solves for a coefficient α in an optimization problem that minimizes the norm of $g(x)$.

State Transition Termination The executor terminates in three different scenarios. (1) If the executor reaches the exit point on the CFG, there will not be any updates to X^{t+1} after Equation 8, naturally leading to an equilibrium state. (2) If the executor reaches an equilibrium state, but not at the program exit point, it logically indicates that further execution will not result in changes in the program state. Therefore, it is natural to terminate. (3) If the executor reaches a configured maximum steps.

Once X^* is at equilibrium, we apply layer normalization [4] and global average pooling layer to obtain the graph representation G :

$$G = \text{LayerNorm}\left(\frac{\sum_i^n X_{i,j}^T}{n}\right), \forall j = 1, \dots, h \quad (11)$$

The prediction task can be simply computed by a linear transformation to get the logits:

$$y' = W_p G, \text{ where } W_p \in \mathbb{R}^{1 \times h} \quad (12)$$

We want to emphasize that through the use of an implicitly defined GNN layer, it is no longer required to have multiple stacking GNN layers to achieve higher order node aggregation. Instead, each state transition within the layer effectively performs a message passing, as a normal GNN layer would. This has the benefits of lowering the memory costs, while maintaining the same level of representational power, given similar parameter count. Moreover, the long range dependency issue can be effectively addressed by iterating a nearly infinite number of state transitions.

4.4 Training

While the forward pass in an implicit network possesses some nice properties for the network discussed earlier, it is not a trivial task to train the backward pass. Traditionally, a neural network contains exact operations with explicitly defined input and output, where the gradients can be computed via chain rule. We first define the loss term l :

$$l = \mathcal{L}(\hat{y}, y) = \mathcal{L}(F_\psi(G), y) \quad (13)$$

F_ψ is the prediction rule that takes the graph embedding G . $\mathcal{L}(\cdot)$ computes the cross entropy loss and outputs the scalar l . Using chain rule, the loss can be backpropagated as:

$$\frac{\partial l}{\partial \theta} = \frac{\partial l}{\partial G} \frac{\partial G}{\partial X^*} \frac{\partial X^*}{\partial \theta} \quad (14)$$

The terms $\frac{\partial l}{\partial G}$ and $\frac{\partial G}{\partial X^*}$ can be both computed using any autograd software. However, the term $\frac{\partial X^*}{\partial \theta}$ is difficult to compute, since the equilibrium point X^* is obtained through iterative root finding. If we unroll this computation graph, the network needs to store all intermediate gradients for every state transition. Depending on the number of transitions, it is not a practical approach. Instead, we write X^* in its implicitly defined form:

$$X^*(\theta) = \phi(X^* W \tilde{A}^* + b_\Omega(U)) = F_I(X^*(\theta), U) \quad (15)$$

where F_I denotes the implicit graph neural network. By taking the derivative with respect to θ , we obtain:

$$\frac{\partial X^*(\theta)}{\partial \theta} = \frac{\partial F_I(X^*(\theta), U)}{\partial \theta} \quad (16)$$

By applying the chain rule on the right hand side of equation 16, we expand it into the following:

$$\frac{\partial X^*(\theta)}{\partial \theta} = \frac{\partial F_I(X^*, U)}{\partial \theta} + \frac{\partial F_I(X^*, U)}{\partial X^*} \frac{\partial X^*(\theta)}{\partial \theta} \quad (17)$$

At this point, both $\frac{\partial F_I(X^*, U)}{\partial \theta}$ and $\frac{\partial F_I(X^*, U)}{\partial X^*}$ can again be obtained using autograd software. The last unknown term $\frac{\partial X^*(\theta)}{\partial \theta}$ is computed by solving the linear system. In our approach, we use Anderson acceleration to iteratively solve this term.

Through implicit differentiation, we directly evaluate the gradient at the equilibrium point. We avoid the computation of any intermediate state transition and can efficiently backpropagate through the network, even with a nearly infinite number of transitions. This also has a better memory footprint.

4.5 Well-posedness

Equation (8) needs to have a unique solution X^* when iterated infinitely. Such property is called the well-posedness. According to Gu et al. [21], W and \tilde{A} are well-posed for ϕ when there is a unique solution. First of all, the choice of ϕ needs to satisfy the component-wise non-expansive (CONE) property, where most activation functions such as ReLU, Sigmoid, and Tanh, possess such property [17]. Then, we need to construct sufficient conditions on W and \tilde{A} with a CONE activation function for well-posedness. It is stated that $\|W\|_\infty < \kappa / \lambda_{pf}(\tilde{A})$ needs to be true, where $\|W\|_\infty$ is the infinity norm, $\lambda_{pf}(\tilde{A})$ is the Perron-Frobenius (PF) eigenvalue [7], $\kappa \in [0, 1)$ is the scaling constant. Equation (8) then has a

unique solution. This is ensured by projecting W in each update to satisfy this condition:

$$W' = \arg \min_{\|M\|_{\infty} \leq \kappa / \lambda_{pf}(\tilde{A})} \|M - W\|_F^2 \quad (18)$$

where $\|\cdot\|_F$ is the Frobenius norm. Note that even with a gated convolution which results in an updated \tilde{A} for every iteration, we still maintain a well-posed \tilde{A} as it contains a strictly smaller or equal PF eigenvalue than the original A , given the agent a is non-expansive, resulting in $\kappa / \lambda_{pf}(\tilde{A}) \geq \kappa / \lambda_{pf}(A)$.

5 EXPERIMENT

In this section, we demonstrate the ability of DeepEXE on predicting binary code vulnerability in a variety of scenarios. To properly evaluate DeepEXE, we conduct experiments using two semi-synthetic datasets and two real world datasets. The NDSS18³ and Juliet Test Suites⁴ are both semi-synthetic datasets commonly used as for vulnerability detection tasks. Though the practical implications for a method should not solely depend on the synthetic results, as they are less complex. For real world datasets that are larger and can contain less trivial vulnerabilities, we employ the FFmpeg⁵ and Esh [12] datasets. For the baseline methods, we inherit the results reported in previous works, due to the large amount of experiments and different setups. The evaluation metrics reported include accuracy, precision, recall, F1 score, and area under the ROC curve (AUC). We randomly split each dataset into 75% for training and 25% for evaluation. Some metrics are not shown in the baselines because of their absence in the original works. The hardware used for the experiments includes a RTX6000 GPU, Intel Xeon Gold 5218 CPU, and 64GB of memory. The main software used includes Python 3.9.10 and PyTorch 1.10.2 on Ubuntu 20.04.3 LTS.

5.1 Datasets

Semi-Synthetic Datasets include the NDSS18 dataset and Juliet Test Suite. The NDSS18 dataset is a derivation from the National Institute of Standards and Technology (NIST): NVD⁶ and the Software Assurance Reference Dataset (SARD) project. NDSS18 was first published by [37] as a source code vulnerability dataset and later compiled to binary code by [32] for binary level detection. It includes a total of 32,281 binary functions that are compiled using Windows and Linux. There are two types of Common Weakness Enumerations (CWEs)⁷ in NDSS18: CWE119 and CWE399. Juliet Test Suite is a collection of 81,000 test cases in C/C++ and Java that contain 112 different CWEs. Both datasets have nearly balanced distributions for the labels.

Real CVE Datasets include the FFmpeg vulnerabilities and Esh datasets, which are both extracted from real world applications or open-source libraries. The codebase is significantly larger than the ones in semi-synthetic datasets. Vulnerabilities are often harder to detect in these programs, due to the much increased complexity. FFmpeg is an open-source suite of libraries written in C for handling media files, such as video and audio. It was first used in source code

vulnerability detection [55], where the authors manually collected and labelled the data for various vulnerability commits on Github. We compile the FFmpeg source code provided by the authors into binary code and obtain 16,494 binary functions, where 7,257 are vulnerable and 9,237 are non-vulnerable. The Esh dataset contains CVE cases collected by David et al. [12], which include 8 different CVEs: cve-2014-0160, cve-2014-6271, cve-2015-3456, cve-2014-9295, cve-2014-7169, cve-2011-0444, cve-2014-4877, and cve-2015-6826. In total, there are 3,379 cases and only 60 are vulnerable. The distribution of vulnerability in the Esh dataset is highly imbalanced, which represents a more realistic scenario.

5.2 Baselines and hyperparameters

NDSS18 baselines Maximal Divergence sequential Autoencoder (MDSAE) was proposed by [32] and uses a deep representation learning approach. The input to MDSAE is a sequence of binary code instructions. We report the three best variants in this paper. MMDSAE [1] is a modified version of MDSAE and uses a similar approach that adds a regularization technique. The authors propose two variants, namely MDSAE-NR and TDNN-NR, that have similar performance. The last baseline we include for this dataset is VulDeePecker [37], which is a source code vulnerability detection. Instead of using binary code instructions, VulDeePecker takes source code gadgets as input. Although this is a source code evaluation, the underlying dataset used is the same.

Juliet baselines Bin2Vec [3] is a graph-based binary vulnerability detection approach that utilizes graph convolution network by taking the CFG as input. Instruction2Vec [34] is a representation learning approach to embed the assembly instructions into vectors and apply the downstream vulnerability detection task. The instruction embedding is similar to Word2Vec, it utilizes different parts of an instruction and combines them as a single vector. The downstream vulnerability detection is achieved by training a CNN or Text-CNN using the vectors. The same authors later proposed an updated version of Instruction2Vec [33] and includes a few variants, including Word2Vec and Binary2Img. All Instruction2Vec related methods use the assembly instructions as input and do not consider the structural information.

FFmpeg baselines For the FFmpeg dataset, we compared DeepEXE to Devign [55], which provides the source code for FFmpeg. Devign uses a gated graph recurrent network (GGRN) [40] as a graph learning technique. Unlike binary code, where only CFG can be extracted, Devign detects vulnerabilities at the source code level. It takes several intermediate graph representations of source code, such as AST, CFG, DFG, and NCS. We include several variants of Devign, such as Bi-LSTM, GGRN with CFG or all graphs, and Devign with CFG or all graphs. Note that we directly compare our results on the binary code with the original results of Devign, which are based on source code.

Esh baselines To the best of our knowledge, Esh is not used in any other papers for vulnerability detection evaluation. The original paper that provided this dataset evaluates it at the basic block level and focuses on code matching. Therefore, we compare DeepEXE to the Bi-LSTM and GCN baselines we implemented ourselves.

For the baseline methods, we directly inherit the results reported in previous works, due to the large amount of experiments and

³<https://samate.nist.gov/SRD/index.php>, Software Assurance Reference Dataset

⁴<https://samate.nist.gov/SARD/test-suites>, NIST Test Suites

⁵<https://ffmpeg.org/>, FFmpeg

⁶<https://nvd.nist.gov/>, National Institute of Standards and Technology

⁷<https://cwe.mitre.org/>, Common Weakness Enumeration (CWE)

Table 1: NDSS18 Dataset Evaluation

Models	Input Type	Accuracy	Recall	Precision	F1	AUC
Bi-LSTM	Assembly Ins.	85.38	83.47	87.09	85.24	94.89
GCN	CFG	86.48	84.59	88.12	86.32	95.81
BERT [15]	Assembly Ins.	50.64	100.00	50.64	67.23	50.00
DistilBERT [42]	Assembly Ins.	85.64	83.01	87.65	85.27	95.42
MD-CWS [32]	Assembly Ins.	85.30	98.10	78.40	87.10	85.20
MD-CKL [32]	Assembly Ins.	82.30	98.00	74.80	84.00	82.10
MD-RWS [32]	Assembly Ins.	83.7	94.3	78.0	85.4	83.5
MDSAE-NR [1]	Assembly Ins.	87.50	99.30	81.20	89.80	87.10
TDNN-NR [1]	Assembly Ins.	86.60	98.70	80.30	88.30	86.30
VulDeePecker [37]	Source Code Gadgets	83.50	91.00	79.50	84.80	83.40
DeepEXE	CFG	90.58	89.36	92.13	90.72	98.01

Table 2: Juliet Dataset Evaluation

Models	Input Type	Accuracy	Recall	Precision	F1	AUC
Bi-LSTM	Assembly Ins.	96.81	98.44	95.48	96.94	99.03
gcn [3]	CFG	97	NA	NA	NA	NA
i2v/CNN [34]	Assembly Ins.	87.6	N/A	N/A	N/A	N/A
i2v/TCNN [34]	Assembly Ins.	96.1	N/A	N/A	N/A	N/A
w2v/CNN [34]	Assembly Ins.	87.9	N/A	N/A	N/A	N/A
w2v/TCNN [34]	Assembly Ins.	94.2	N/A	N/A	N/A	N/A
i2v [33]	Assembly Ins.	96.81	97.07	96.65	96.85	N/A
bin2img [33]	Assembly Ins.	97.53	97.05	97.91	97.47	N/A
w2v [33]	Assembly Ins.	96.01	96.07	95.92	95.99	N/A
DeepEXE	CFG	99.80	99.60	100.00	99.80	100.00

Table 3: FFmpeg Dataset Evaluation

Models	Code Level	Input Type	Accuracy	F1
Bi-LSTM [55]	Source Code	Code Snippets	53.27	69.51
Bi-LSTM + Attention [55]	Source Code	Code Snippets	61.71	66.01
CNN [55]	Source Code	Code Snippets	53.42	66.58
GGRN-CFG [55]	Source Code	CFG	65.00	71.79
GGRN-composite [55]	Source Code	AST, CFG, DFP, NCS	64.46	70.33
Devign-CFG [55]	Source Code	CFG	66.89	70.22
Devign-composite [55]	Source Code	AST, CFG, DFP, NCS	69.58	73.55
DeepEXE	Binary Code	CFG	68.29	67.17

Table 4: Esh Dataset Evaluation

Models	Input Type	Accuracy	Recall	Precision	F1	AUC
Bi-LSTM	Assembly Ins.	99.49	79.48	88.57	83.78	96.87
GCN	CFG	99.31	63.89	95.83	76.67	83.54
DeepEXE	CFG	99.78	95.65	91.67	93.62	99.78

different setups. The evaluation metrics reported include accuracy, precision, recall, F1 score, and area under the ROC curve (AUC). Cross-validation is used for tuning hyperparameters in order to obtain optimal accuracy and reasonable memory usage. We use a universal hidden dimension of 64, learning rate of 0.01 with the Adam optimizer [28], dropout rate of 0.5, batch size of 192, and

a maximum iteration of 50 for the Anderson acceleration solver. We randomly split each dataset into 75% for training and 25% for evaluation. Some metrics are not shown in the baselines because of their absence in the original works. The hardware used for the experiments includes a RTX6000 GPU, Intel Xeon Gold 5218 CPU,

and 64GB of memory. The main software used includes Python 3.9.10 and PyTorch 1.10.2 on Ubuntu 20.04.3 LTS.

5.3 Evaluation

For each dataset, we compare DeepEXE with the benchmarks that are also evaluated on the same dataset due to limited space. Additionally, we build baseline models that work inherently well in this task including bi-directional LSTM (Bi-LSTM) [25] and graph convolution network (GCN) [29].

Semi-Synthetic Results We first analyze the results for the NDSS18 dataset shown in Table. 1. As the popularity of Large Language Models (LLMs) grows drastically in the recent years, we briefly showcase two pre-trained LLMs, DistilBERT [42] and BERT [15], to perform the downstream vulnerability detection tasks for the NDSS dataset. Several baselines methods including Bi-LSTM, GCN, and DistilBERT compare relatively well with the benchmark. However, the original BERT model struggles due to the tokenization method. As a result, LLMs do not provide significant improvements compared to RNNs in our case. All the other benchmarks have imbalanced precision and recall, where the models tend to overestimate the vulnerable code. DeepEXE has the best overall performance, leading the accuracy and AUC by 3%. Moreover, DeepEXE is a CFG-based method and we empirically show that by adding the execution-guided agent and expanding the receptive field of graph convolution, it is able to capture more topological information. DeepEXE also outperforms VulDeePecker, which is a source code level method that only leverages the sequential information of the code gadget, potentially omitting the much useful topological knowledge of the source code.

The Juliet dataset evaluation is shown in Table 2. As a synthetic dataset, the test cases contain much shorter code. However, there are over 100 different CWEs among all test cases. In reality, a detection tool should be robust enough to detect unseen or zero-day vulnerabilities. It is useful for evaluating the robustness and generalizability of an approach. DeepEXE shows nearly perfect detection accuracy and AUC for this dataset. This shows that even with the single-layer design, DeepEXE is able to generalize well enough. As the graphs are usually small in these test cases, the execution paths generated by static analysis are likely more accurate. Therefore, we believe the implicit GNN contributes more to the performance increase than the agent in this case.

Real CVE Results We evaluate the FFmpeg dataset shown in Table 3, which specifies the code levels and input types. Since Devign detects vulnerabilities at the source code level, it is significantly easier with the rich semantics, syntax, and structures. DeepEXE is able to outperform most of the approaches, even at the binary code level. In particular, when only using the CFG as input, DeepEXE achieves better accuracy than both the Devign and GGRN models. Devign-composite utilizes multiple input graphs, such as AST, DFP and NCS. These additional graphs are usually only available for source code. DeepEXE shows its capability at detecting vulnerabilities for real-world and complex programs. Moreover, source code CFGs are less complicated to generate, whereas binary CFGs often can be an over-estimation of the true control flow. With our execution-guided approach, we limit the errors caused by such approximation, while maintaining a high level of global information.

The receptive field of GNN in DeepEXE is practically unlimited, allowing us to accommodate for much larger graphs.

Lastly, we show the evaluation results for the Esh dataset in Table 4. Due to the extreme imbalance of labels distribution, which is the case in many real-life scenarios, the Bi-LSTM and GCN baselines have lower recalls. The recall metric is important when there are fewer vulnerable cases. DeepEXE, on the other hand, is able to distinguish vulnerable code from non-vulnerable code, given the small number of positive labels. Note that the class weight is not manually adjusted during training, as it is cumbersome and inefficient to tune it for every dataset in practice. With over 90% precision, DeepEXE is able to identify 95% of the vulnerable CVE cases. Similar to FFmpeg, although many cases in the Esh dataset contain a large number of nodes, DeepEXE is inherently designed to handle such large graphs and outperform other baselines.

6 CONCLUSIONS

We have proposed DeepEXE, a control flow execution-guided deep learning framework for binary code vulnerability detection. Given the importance of binary code learning, we address two major gaps in the existing research works, which are the lack of modelling program state transition and scalability for large graphs. Instead of assuming the CFG is accurate, which is often not the case, due to the over-estimation from static analysis, we use a reinforcement agent to guide the execution of a program flow that mimics the behaviour of dynamic analysis. DeepEXE is able to capture certain program state transitions that lead to specific vulnerability results, creating a higher dependency between the output and internal node state and topological information. We also show the benefits of training an implicitly defined network, which are directly obtaining the gradients for the equilibrium point and mitigating the heavy memory footprint in large networks. In the experiments, we demonstrate that DeepEXE outperforms all state-of-the-art vulnerability detection methods for the NDSS18 and Juliet datasets. DeepEXE is also very competitive in detecting real world CVEs, even when compared to source code level methods, which are less difficult, given the amount of available information. Overall, DeepEXE is a robust and accurate tool for binary vulnerability detection.

In the future, there are several potential directions to grow for DeepEXE. First of all, the training time is slower than for the traditional neural network, due to the many iterations for obtaining equilibrium. This can be improved by using more sophisticated solvers to reduce the number of steps for equilibrium computation. Next, DeepEXE does not have to be restricted to vulnerability detection in the cybersecurity domain. For other security tasks, such as binary code similarity comparison or malware detection, matching the graph structures of malicious programs is often done using GNN. By modifying the training objective, DeepEXE can be used for a lot more of supervised and unsupervised tasks. Moreover, as long as the input data has some form of graphical structures, we can apply the same design to many other domains, such as social network and chemistry studies.

7 ACKNOWLEDGEMENT

This research is supported by BlackBerry Limited. (ALLRP 561035) and NSERC Alliance Grants (ALLRP 561035-20).

REFERENCES

- [1] Marwan Ali Albahar. 2020. A Modified Maximal Divergence Sequential Auto-Encoder and Time Delay Neural Network Models for Vulnerable Binary Codes Detection. *IEEE Access* 8 (2020), 14999–15006.
- [2] Nikolaos Alexopoulos, Sheikh Mahbub Habib, Steffen Schulz, and Max Mühlhäuser. 2020. The tip of the iceberg: On the merits of finding security bugs. *ACM Transactions on Privacy and Security (TOPS)* 24, 1 (2020), 1–33.
- [3] Shushan Arakelyan, Christophe Hauser, Erik Kline, and Aram Galstyan. 2020. Towards Learning Representations of Binary Executable Files for Security Tasks. *arXiv preprint arXiv:2002.03388* (2020).
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [5] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. 2019. Deep equilibrium models. *Advances in Neural Information Processing Systems* 32 (2019).
- [6] Roberto Baldoni, Emilio Coppa, Daniele Cono D'elia, Camil Demetrescu, and Irene Finocchi. 2018. A survey of symbolic execution techniques. *ACM Computing Surveys (CSUR)* 51, 3 (2018), 1–39.
- [7] Abraham Berman and Robert J Plemmons. 1994. *Nonnegative matrices in the mathematical sciences*. SIAM.
- [8] Sicong Cao, Xiaobing Sun, Lili Bo, Ying Wei, and Bin Li. 2021. Bgmn4vd: constructing bidirectional graph neural-network for vulnerability detection. *Information and Software Technology* 136 (2021), 106576.
- [9] Mahinthan Chandramohan, Yinxing Xue, Zhengzi Xu, Yang Liu, Chia Yuan Cho, and Hee Beng Kuan Tan. 2016. BinGo: cross-architecture cross-OS binary search. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*.
- [10] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).
- [11] Hanjun Dai, Zornitsa Kozareva, Bo Dai, Alex Smola, and Le Song. 2018. Learning steady-states of iterative algorithms over graphs. In *International conference on machine learning*. PMLR, 1106–1114.
- [12] Yaniv David, Nimrod Partush, and Eran Yahav. 2016. Statistical similarity of binaries. *Acm Sigplan Notices* 51, 6 (2016), 266–280.
- [13] Yaniv David and Eran Yahav. 2014. Tracelet-based code search in executables. *Acm Sigplan Notices* 49, 6 (2014), 349–360.
- [14] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. 2018. End-to-end differentiable physics for learning and control. *Advances in neural information processing systems* 31 (2018).
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [16] Yue Duan, Xuezixiang Li, Jinghan Wang, and Heng Yin. 2020. Deepbindiff: Learning program-wide code representations for binary diffing. In *Network and Distributed System Security Symposium*.
- [17] Laurent El Ghaoui, Fangda Gu, Bertrand Travacca, Armin Askari, and Alicia Tsai. 2021. Implicit deep learning. *SIAM Journal on Mathematics of Data Science* 3, 3 (2021), 930–958.
- [18] Kathryn A Farris, Ankit Shah, George Cybenko, Rajesh Ganesan, and Sushil Jajodia. 2018. Vulcon: A system for vulnerability prioritization, mitigation, and management. *ACM Transactions on Privacy and Security (TOPS)* 21, 4 (2018), 1–28.
- [19] Claudio Gallicchio and Alessio Micheli. 2010. Graph echo state networks. In *The 2010 international joint conference on neural networks (IJCNN)*. IEEE, 1–8.
- [20] 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*, Vol. 2. 729–734.
- [21] Fangda Gu, Heng Chang, Wenwu Zhu, Somayeh Sojoudi, and Laurent El Ghaoui. 2020. Implicit graph neural networks. *Advances in Neural Information Processing Systems* 33 (2020), 11984–11995.
- [22] Aakanshi Gupta, Bharti Suri, Vijay Kumar, and Pragyashree Jain. 2021. Extracting rules for vulnerabilities detection with static metrics using machine learning. *International Journal of System Assurance Engineering and Management* 12, 1 (2021), 65–76.
- [23] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *arXiv preprint arXiv:1706.02216* (2017).
- [24] Jacob A Harer, Louis Y Kim, Rebecca L Russell, Onur Ozdemir, Leonard R Kosta, Akshay Rangamani, Lei H Hamilton, Gabriel I Centeno, Jonathan R Key, Paul M Ellingwood, et al. 2018. Automated software vulnerability detection with machine learning. *arXiv preprint arXiv:1803.04497* (2018).
- [25] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [26] Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144* (2016).
- [27] James C King. 1976. Symbolic execution and program testing. *Commun. ACM* 19, 7 (1976), 385–394.
- [28] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014). <https://api.semanticscholar.org/CorpusID:6628106>
- [29] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [30] Taku Kudo. 2018. Subword regularization: Improving neural network translation models with multiple subword candidates. *arXiv preprint arXiv:1804.10959* (2018).
- [31] Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226* (2018).
- [32] Tue Le, Tuan Nguyen, Trung Le, Dinh Phung, Paul Montague, Olivier De Vel, and Lizhen Qu. 2018. Maximal divergence sequential autoencoder for binary software vulnerability detection. In *International Conference on Learning Representations*.
- [33] Yongjun Lee, Hyun Kwon, Sang-Hoon Choi, Seung-Ho Lim, Sung Hoon Baek, and Ki-Woong Park. 2019. Instruction2vec: Efficient Preprocessor of Assembly Code to Detect Software Weakness with CNN. *Applied Sciences* 9, 19 (2019), 4086.
- [34] Young Jun Lee, Sang-Hoon Choi, Chulwoo Kim, Seung-Ho Lim, and Ki-Woong Park. 2017. Learning binary code with deep learning to detect software weakness. In *KSII The 9th International Conference on Internet (ICONI) 2017 Symposium*.
- [35] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. 2015. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493* (2015).
- [36] Zhen Li, Deqing Zou, Shouhuai Xu, Zhaoxuan Chen, Yawei Zhu, and Hai Jin. 2021. Vuldelocator: a deep learning-based fine-grained vulnerability detector. *IEEE Transactions on Dependable and Secure Computing* (2021).
- [37] Zhen Li, Deqing Zou, Shouhuai Xu, Xinyu Ou, Hai Jin, Sujuan Wang, Zhijun Deng, and Yuyi Zhong. 2018. VulDeePecker: A deep learning-based system for vulnerability detection. *arXiv preprint arXiv:1801.01681* (2018).
- [38] Bingchang Liu, Wei Huo, Chao Zhang, Wenchao Li, Feng Li, Aihua Piao, and Wei Zou. 2018. α diff: cross-version binary code similarity detection with dnn. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. 667–678.
- [39] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [40] Luana Ruiz, Fernando Gama, and Alejandro Ribeiro. 2020. Gated graph recurrent neural networks. *IEEE Transactions on Signal Processing* 68 (2020), 6303–6318.
- [41] Sefa Eren Şahin, Ecem Mine Özyedierler, and Ayse Tosun. 2022. Predicting vulnerability inducing function versions using node embeddings and graph neural networks. *Information and Software Technology* 145 (2022), 106822.
- [42] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
- [43] 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.
- [44] Kazi Zakia Sultana, Vaibhav Anu, and Tai-Yin Chong. 2021. Using software metrics for predicting vulnerable classes and methods in Java projects: A machine learning approach. *Journal of Software: Evolution and Process* 33, 3 (2021), e2303.
- [45] Donghai Tian, Xiaoqi Jia, Rui Ma, Shuke Liu, Wenjing Liu, and Changzhen Hu. 2021. BinDeep: A deep learning approach to binary code similarity detection. *Expert Systems with Applications* 168 (2021), 114348.
- [46] Junfeng Tian, Wenjing Xing, and Zhen Li. 2020. BVDetector: A program slice-based binary code vulnerability intelligent detection system. *Information and Software Technology* 123 (2020), 106289.
- [47] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
- [48] Homer F Walker and Peng Ni. 2011. Anderson acceleration for fixed-point iterations. *SIAM J. Numer. Anal.* 49, 4 (2011), 1715–1735.
- [49] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3 (1992), 229–256.
- [50] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [51] Tao Xie, Nikolai Tillmann, Jonathan De Halleux, and Wolfram Schulte. 2009. Fitness-guided path exploration in dynamic symbolic execution. In *2009 IEEE/IFIP International Conference on Dependable Systems & Networks*. IEEE, 359–368.
- [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] Han Yan, Senlin Luo, Limin Pan, and Yifei Zhang. 2021. HAN-BSVD: a hierarchical attention network for binary software vulnerability detection. *Computers & Security* 108 (2021), 102286.
- [54] Zeping Yu, Rui Cao, Qiyi Tang, Sen Nie, Junzhou Huang, and Shi Wu. 2020. Order matters: semantic-aware neural networks for binary code similarity detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1145–1152.

- [55] Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. 2019. Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks. In *Advances in Neural Information Processing Systems*. 10197–10207.
- [56] Deqing Zou, Sujuan Wang, Shouhuai Xu, Zhen Li, and Hai Jin. 2019. μ VulDeePecker: A Deep Learning-Based System for Multiclass Vulnerability Detection. *IEEE Transactions on Dependable and Secure Computing* 18, 5 (2019), 2224–2236.
- [57] Fei Zuo, Xiaopeng Li, Patrick Young, Lannan Luo, Qiang Zeng, and Zhixin Zhang. 2018. Neural machine translation inspired binary code similarity comparison beyond function pairs. *arXiv preprint arXiv:1808.04706* (2018).