

Learning to Reverse DNNs from AI Programs Automatically

Simin Chen¹, Hamed Khanpour², Cong Liu¹ and Wei Yang¹

¹The University of Texas at Dallas

²Microsoft Research

{Simin.Chen, Cong, Wei.Yang}@utdallas.com, Hamed.Khanpour@microsoft.com,

Abstract

With the privatization deployment of DNNs on edge devices, the security of on-device DNNs has raised significant concern. To quantify the model leakage risk of on-device DNNs automatically, we propose NNReverse, the first learning-based method which can reverse DNNs from AI programs without domain knowledge. NNReverse trains a representation model to represent the semantics of binary code for DNN layers. By searching the most similar function in our database, NNReverse infers the layer type of a given function’s binary code. To represent assembly instructions semantics precisely, NNReverse proposes a more fine-grained embedding model to represent the textual and structural-semantic of assembly functions.

1 Introduction

Over the last few years, deploying Deep Neural Networks (DNNs) on edge devices is becoming a popular trend for giant AI providers. The AI providers compile the private high-quality DNN models into stand-alone programs [Li *et al.*, 2020; Chen *et al.*, 2018a; Rotem *et al.*, 2018; Cyphers *et al.*, 2018] and would like to sell them to other companies, organizations, and governments with a license fee. This edge-deployment situation further exacerbates the risk of model leakage. The DNN model leakage will result in the following three consequences. First, the potential loss of intellectual property. DNNs parameters are typically trained from TB datasets with high training costs. For example, training a DNN on the google cloud TPU will cost more than \$400K [Devlin *et al.*, 2018; Raffel *et al.*, 2019; Zhu *et al.*, 2021]. Thus, DNNs architecture and parameters are precious intellectual property of the developers. Second, the potential infringement of privacy. DNNs parameters have been proven to remember sensitive information (*e.g.* SSN) in the DNNs training data [Carlini *et al.*, 2020]. Therefore, the leakage of DNN parameters will result in the privacy leakage of training data. Last but not least, the potential risk of being attacked by adversaries. DNNs are vulnerable to adversarial attacks [Carlini and Wagner, 2017]. If the attackers obtain the parameters, they could easily compute DNNs

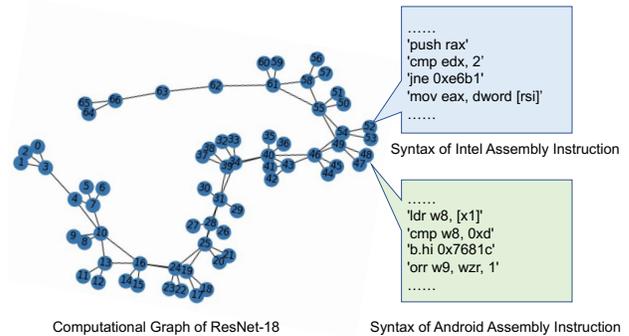


Figure 1: Assembly instructions of ResNet-18 on different platforms

gradients to launch white-box attacks. For Example, reversing Apple’s image hash model, NeuralHash ¹, will enable attackers to break the content protection provided by macOS. Therefore, conducting penetration testing to protect DNNs’ privacy is essential.

Quantifying model leakage risk of on-device DNNs automatically is a challenging task because of the following two reasons: (1) Reversing DNN architectures requires understanding the semantics of the machine code of each DNN layer. Machine code is incomprehensible to humans; unfortunately, modern DNN usually contains hundreds of layers, and each layer corresponds to a function that needs to infer the semantic. (2) The cross-platform deployment of DNN models increase the difficulty of understanding the semantics of the code. DNNs may be deployed on different hardware platforms. The machine code on different hardware platforms have different syntax. Therefore, understanding the semantics of machine code on different hardware platforms requires different domain-specific knowledge. Fig. 1 shows the challenge in reversing DNNs, the left figure shows the computational graph of Resnet-18, reversing ResNet-18 requires inferring the semantic of each node in the graph, and the right figure shows different assembly instructions syntax on different hardware platforms.

In this paper, we propose NNReverse, the first static model stealing attack to reverse DNNs from AI programs automatically. Different from existing model extraction at-

¹Apples-neuralhash-algorithm-has-been-reverse-engineered

tacks [Jagielski *et al.*, 2020; Oh *et al.*, 2019; Tramèr *et al.*, 2016], which approximate the victim DNNs functionality rather than reversing DNN architectures [Krishna *et al.*, 2019]. Our approach does not train a substitute model to approximate input-output pairs from the victim DNNs. Instead, our approach learns and infer the semantic of machine code for different hardware platforms automatically.

Our intuition is that although DNN architectures are complicated, the basic types (*i.e.* *conv*, *dense*) of DNN layers are enumerable. Thus, it is possible to reverse DNN architectures by inferring each DNN layer type. Furthermore, because different DNN layers are implemented as different binary functions, we can infer DNN layer type by inferring semantic corresponding binary functions. Specifically, our attack includes two phases: offline phase and online phase. In the offline phase, we build a database to collect the binary of each optimized kernel layer, then we train a semantic representation model on the decompiled functions. The semantic representation model can be trained with machine code from different hardware platforms, thus, NNReverse could reverse DNNs for different hardware platforms. Although some semantic representation approaches for NLP fields have been proposed, they are not suitable for assembly functions because they ignore the topology semantic in assembly functions. We propose a new embedding algorithm to overcome this limitation and represent assembly functions more precisely, which combines syntax semantic representation and topology semantic. In the online phase, we decompile the victim DNN executable and collect the binaries. By the pre-collected knowledge in the database, we reconstruct the whole DNN model. We summarize our contribution as follows:

- **Characterizing:** We characterize the model leakage risk of on-device DNNs. Specifically, we show that on-device DNNs could be reversed easily.
- **Approach:** We propose a new fine-grained embedding approach for representing assembly functions, which combines syntax representation and topology structure representation.
- **Evaluation:** We implement NNReverse and evaluate NNReverse on two datasets. The evaluation results show NNReverse can achieve higher accuracy in searching similar semantic functions than baseline methods, and we apply NNReverse to reverse real-world DNNs as a case study. Results show that NNReverse can reverse the DNNs without accuracy loss.
- **Dataset:** We open-source our decompiled assembly function dataset², including more than 80k functions compiled for four hardware platforms. Our dataset is the first dataset for DNN reverse engineering, contributing to future research on on-device DNN security.

2 Background & Related Work

2.1 DNN Deployment

DL compilers are used to generate portable AI programs from high-level DL frameworks such as PyTorch,

²<https://drive.google.com/file/d/NNReverse>

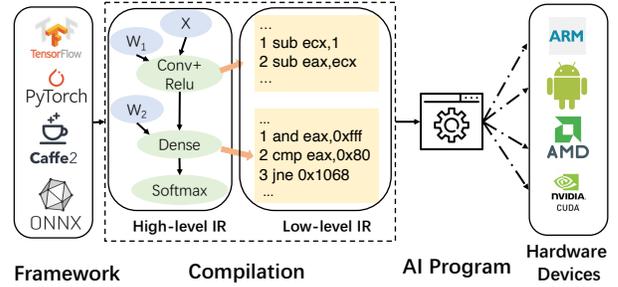


Figure 2: The process of compiling DNNs into programs.

Work	Information	Method	Platform-Independent	Results	
				Architecture	Parameters
1 [Yan <i>et al.</i> , 2020]	Catche	Search	N	×	✓
2 [Hua <i>et al.</i> , 2018]	Accelerator	Search,Infer	Y	✓	×
3 [Duddu <i>et al.</i> , 2018]	Timing	Search	N	✓	×
4 [Roberts <i>et al.</i> , 2019]	Noise Input	Infer	Y	P	P
5 [Jagielski <i>et al.</i> , 2020]	Queries	Infer	Y	×	×
Ours	Binary files	Infer	Y	✓	✓

Table 1: Y stands for applicable for specific device, N stands for device-independent. ✓ stands for fully recover, P stands for partial recover, × means cannot recover.

TensorFlow, and Caffe. The generated AI programs are advanced optimized (*e.g.* Layer Fusion) to match the target hardware devices. The working process of DL compiler is shown in Fig. 2. DL compilers accept DL models from high-level DL frameworks as input and first translate DL models into high-level intermediate representation (IR). The high-level IR (as known as graph IR) is a computational graph, where nodes represent the atomic DL operators (Convolution, Pooling, Convolution+ReLU, *etc.*) or tensors, and the edges represent the direction of dataflow. Then DL compilers translate each atomic DL operator into low-level IR and generate the executable code by optimizing the computation and memory access for the target hardware devices. DL compilers enable running DL models on different hardware devices such as X86, ARM, NVIDIA, and TPU. For more details about DL compilers, we refer readers to the existing work [Li *et al.*, 2020; Chen *et al.*, 2018a; Rotem *et al.*, 2018; Cyphers *et al.*, 2018]. In this paper, we consider TVM [Chen *et al.*, 2018b], a widely-used DL compiler to deploy DNNs. Given a DNNs from high-level frameworks (*e.g.* PyTorch), TVM will output three files: graph file, parameter file, and binary file. Notice that, These files are necessary for compiled AI programs to run. There are similar functionality files for other DL compilers. For example, Apple Core ML applies the file formats *.espresso.net* and *.espresso.weights* to store computational graphs and weights.

2.2 Model Extraction Attacks

As shown in Table 1, a number of attack algorithms [Jagielski *et al.*, 2020; Oh *et al.*, 2019; Tramèr *et al.*, 2016; Zhu *et al.*, 2021] have been proposed to extract DNNs. We summarize the existing attacks into two categories based on their extraction fidelity. The first type of existing attack is *functionality attack*, they treat the victim model as a label or-

acle and train another model to approximate the input-output pair from the victim model. These attacks can not steal the DNN architecture and the DNN parameters, instead, they just clone the functionality of the victim model. The second type of existing attack is *steal attack*. The goal of this attack is to reverse the actual DNN architecture or parameters. For example, [Yan *et al.*, 2020] introduced a cache-based side-channel attack to steal DNN architectures, [Roberts *et al.*, 2019] assume the model architecture and the output of the softmax layer is known, and replicate the model parameters by feeding noise inputs. Different from the existing *steal attack*, we propose an infer-based *steal attack* which requires no assumption about the environment in which the DNN model is running.

2.3 Binary Code Analysis

In this section, we introduce some binary analysis techniques for reverse engineering. Traditional binary analysis techniques often extract hand-crafted features by domain experts to match similar functions and infer binary function semantics. Recently, motivated by the representation learning in NLP fields, researchers propose to apply learning-based approaches for binary code analysis. Specifically, researchers propose to learn a function representation that is supposed to encode the binary function syntax and semantics in low dimensional vectors, known as function embeddings. For example, inspired by Doc2Vec, Asm2Vec proposes to model binary code as documents and design a novel learning algorithm to embed binary functions. However, the learning architectures adopted in existing approaches are limited in modeling long-range dependencies (based on Doc2Vec). The binary code for the DNN operators is long-range dependent because of the loop operations. For example, the four-level loops in the Conv2D layer. Unlike existing approaches, our approach models both instruction and topology semantics of binary code, thus representing the semantics more precisely.

3 NNReverse

3.1 Problem Formulation

In this paper, we consider an AI model privatization deployment scenario, where the DNN model developers compile their private DNN models into stand-alone AI programs and sell the programs to third-party customers with subscription or perpetual licensing. The considered scenario is realistic because many AI models have been deployed on mobile devices (*e.g.*, Apple NeuralHash, Mobile Translation App, AIoT) and this assumption is widely used in existing work. The end-users can physically access the AI programs, especially the machine code (executable code) of the AI programs. The goal of malicious attackers is to reverse the DNNs from the AI program. Specifically, they can reverse the architecture of the DNNs and reconstruct the DNNs with high-level frameworks (*i.e.* Pytorch) to compute the DNNs gradients to launch evasion attacks. Formally, our goal includes: (i) inferring the kernel operator type of each binary function in the victim AI-programs; (ii) reconstructing whole DNNs with high-level DL frameworks. Reversing on-device DNNs have a severe real-world impact because DNNs parameters are precious intellectual property, and reverse engineering can bridge the gap

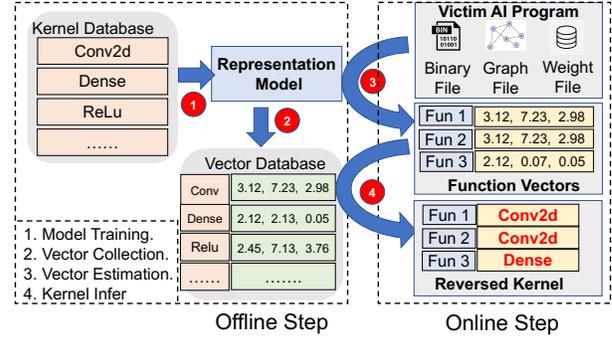


Figure 3: Design overview of our attack method

between white-box and black-box attacks for DNNs (*e.g.*, bypass Apple NeuralHash). Furthermore, understanding the semantic intricacies of deployed binary code makes automatically reversing on-device DNNs difficult.

3.2 Design Overview

As shown in Fig. 3, our approach includes offline and online steps. In the offline step, we collect a set of DNNs with known architectures and use TVM [Chen *et al.*, 2018b], a widely-used DL compiler, to compile them into AI programs including a binary file for each program. We keep the debugging information for each binary file to keep a record of the kernel operator type of each binary function. We then train a semantic-representation model with these binary files. In the online step, given a victim AI program, we first parse the AI program and decompile the parsed stripped binary file into a list of assembly functions. Next, we feed each decompiled assembly function into the trained semantic representation model to estimate the vectors. Finally, we compare the estimated vectors with the vectors in our offline database to infer the kernel operator type. We could reverse the DNNs arch after inferring the kernel operator type and rebuild it with PyTorch by combing the weight files.

3.3 Semantic Representation Model

In this section, we propose an embedding model to represent assembly functions as continuously distributed vectors. As shown in Fig. 4, an assembly function can be represented as a control flow graph (CFG), where nodes are blocks consisting of sequential assembly instructions and edges are control flow paths. Our embedding paradigm, in particular, fuses both text representation and structure representation from the assembly functions.

Text Embedding. We first model the assembly functions as sequential assembly instructions and seek to learn a function $\phi_t(\cdot)$, whose input is a sequential assembly instruction and output is a N dimensional vector. We embed the semantics of assembly instructions with more fine-grained modeling rather than treating each instruction as a token. Specifically, for each assembly instruction, we differentiate the *opcode* and *operands* in the instruction. Then we concatenate the vector of *opcode* and *operands* as the representation of the instruction. For instruction with more than one *operands*, we use the

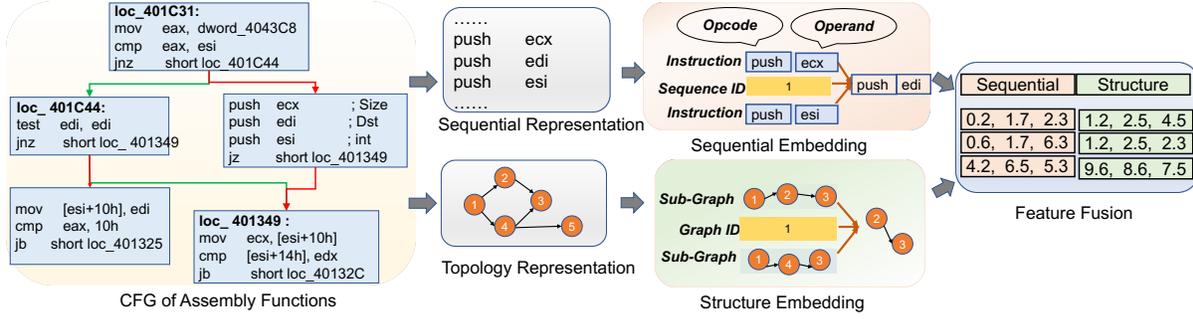


Figure 4: Design overview of our semantic representation model

average vector instead. It can be formulated as:

$$\mathcal{I}(ins) = \mathbb{E}(\Omega(ins)) \parallel \frac{1}{|\Theta(ins)|} \sum_{i=0}^{|\Theta(ins)|} \mathbb{E}(\Theta(ins)_i) \quad (1)$$

Where ins represents an assembly instruction, $\Omega(\cdot)$ and $\Theta(\cdot)$ are the functions to parse *opcode* and *operands* respectively, and $\mathbb{E}(\cdot)$ is the basic embedding layer that transfer tokens to corresponding numeric vectors. Take the assembly instruction $ins = push, rdx$ as example, then $\Omega(ins) = push$ and $\Theta(ins) = rdx$. Because the *operands* may be a memory address (e.g. $0x73ff$) or a immediate value (e.g. $0x16, 0x1024$). To eliminate the issue of out of vocabulary (OOV), we propose a blur strategy. Specifically, for memory address value and immediate value, we compare the value with 1000. If the value is less than 1000, we keep the original value, otherwise, we use token *large* to replace the original value.

$$\sum_{j=1}^M \sum_{i=1}^{|\mathcal{T}(f_j)|} P(\mathcal{I}_i(ins) | C_i, \phi_t(f_j)) \quad (2)$$

After modeling assembly instruction, we seek to represent the semantic of the assembly function. Specifically, we treat each instruction $\mathcal{I}(ins)$ as a “word” and apply a similar idea with doc2vec. Formally, our training objective can be formulated as Equation 2. Where M is the number of the training assembly functions, f_j is the j^{th} assembly function, $T(\cdot)$ is the text representation, C_i is the neighbor instruction of $\mathcal{I}_i(ins)$ and ϕ_t is the text embedding under learned.

Algorithm 1 Structure Embedding Algorithm

Input: G . {A set of CFGs topology structures}
Input: D . {Maximum degree of rooted subgraphs}
 1: Initialization: Sample ϕ_s from $R^{G \times N}$
 2: **for** $g_i \in G$ **do**
 3: **for** $n \in \text{NODE}(g_i)$ **do**
 4: **for** $d = 0$ to D **do**
 5: $sub_g = \text{ComputeWLSubgraph}(n, g_i, d)$
 6: $\phi_s = \phi_s - \frac{\partial -\log P(sub_g | \phi_s(g_i))}{\partial \phi_s}$
 7: **end for**
 8: **end for**
 9: **end for**

Structure Embedding. Given a set of CFG topology structures, we intend to learn a distributed representation for every CFGs’ topology structure. Formally, we represent each CFG as a graph $G = (N, E)$, where N are assembly blocks and E are connections between blocks. Our goal is to learn a embedding function $\phi_s(\cdot)$, whose input is $g \in G$ and output is a \mathbb{N} dimensional vector. Our intuition is that two graphs are more topology similar if these two graphs have more topology similar sub-graphs. Different from the semantic embedding that we apply random-walk paths as representation, applying sub-graphs will keep higher-order substructure and non-linear substructure information. As shown in Algorithm 1. Our algorithm require a set of graphs G as the training corpus. In each training iteration, we compute the WeisfeilerLehman sub-graphs [Shervashidze *et al.*, 2011] of each node in each graph g_i , then we try to maximize the possibility that the sub-graphs appear under the whole graph and update our structure embedding function $\phi_s(\cdot)$. We repeat this process iteratively until the maximum epochs is reached.

$$\phi_{binary}(f) = \frac{\phi_t(T(f))}{|\phi_t(T(f))|} \parallel \frac{\phi_s(S(f))}{|\phi_s(S(f))|} \quad (3)$$

Feature Fusion. Finally, we combine the vectors from text embedding and structure embedding to obtain the final function embedding. Equation 3 shows our fusion process, where f is the assembly function, $T(f)$ and $S(f)$ are the text representation and the structure representation of f respectively, ϕ_t and ϕ_s are the text embedding function and structure embedding function, and \parallel is the vector concatenation.

$$kernel(\mathcal{B}) = name_i \quad i = \text{argmax}_j \frac{v \cdot u_j}{\|v\| \times \|u_j\|} \quad (4)$$

Binary Matching. After the model is well trained, we obtain the pre-trained database $U = \{(u_j, name_j)\}$, where u_j and $name_j$ are the vector and the kernel type of the j^{th} function in the database. During the inference step, given a binary function \mathcal{B} , we first feed \mathcal{B} into our model to get a vector v . we then apply Eq.(4) to assign \mathcal{B} a known kernel type.

3.4 DNN Reconstruction

After training a semantic representation model ϕ_{binary} . We propose a automatic algorithm to rebuild the victim DNNs. As shown in Algorithm 2, the inputs of our algorithm are

Dataset		P			R			F1		
		Text	Structure	Combine	Text	Structure	Combine	Text	Structure	Combine
Vision	ARM	0.8595	0.5018	0.8663	0.8875	0.6665	0.8855	0.8733	0.5725	0.8758
	Aarch64	0.7562	0.3468	0.7688	0.8406	0.5388	0.8306	0.7962	0.4220	0.7985
	Android	0.8419	0.4578	0.8398	0.8741	0.6746	0.8766	0.8577	0.5455	0.8578
	Intel	0.8427	0.4954	0.8444	0.8725	0.6719	0.8713	0.8574	0.5703	0.8576
	AVG	0.8251	0.4505	0.8298	0.8687	0.6380	0.8660	0.8462	0.5276	0.8474
Textual	ARM	0.7232	0.2337	0.7610	0.6859	0.3990	0.7104	0.7041	0.2948	0.7348
	Aarch64	0.7301	0.5148	0.7342	0.6224	0.2566	0.6519	0.6720	0.3425	0.6906
	Android	0.6836	0.2007	0.6660	0.6872	0.2847	0.6882	0.6854	0.2354	0.6769
	Intel	0.7557	0.4085	0.7146	0.6889	0.3825	0.7050	0.7208	0.3951	0.7098
	AVG	0.7232	0.3394	0.7190	0.6711	0.3307	0.6889	0.6956	0.3170	0.7030

Table 2: The results of proposed text embedding, structure embedding and combined embedding

Data	Platform	word2vec (skip)			word2vec (cbow)			doc2vec (skip)			doc2vec (cbow)			Asm2Vec			Ours		
		P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Vision	ARM	0.44	0.48	0.46	0.43	0.49	0.46	0.57	0.71	0.63	0.63	0.72	0.68	0.62	0.76	0.68	0.87	0.89	0.88
	Aarch64	0.29	0.41	0.34	0.45	0.37	0.40	0.43	0.60	0.50	0.43	0.61	0.50	0.59	0.66	0.62	0.77	0.83	0.80
	Android	0.42	0.50	0.46	0.41	0.51	0.46	0.57	0.71	0.63	0.61	0.71	0.66	0.64	0.65	0.64	0.84	0.88	0.86
	Intel	0.42	0.48	0.45	0.43	0.48	0.45	0.57	0.68	0.62	0.58	0.69	0.63	0.62	0.73	0.67	0.84	0.87	0.86
	mixed	0.39	0.47	0.42	0.43	0.46	0.44	0.53	0.68	0.60	0.56	0.68	0.62	0.62	0.70	0.65	0.83	0.87	0.85
Textual	ARM	0.41	0.41	0.41	0.38	0.36	0.37	0.51	0.44	0.48	0.47	0.43	0.45	0.58	0.59	0.59	0.76	0.71	0.73
	Aarch64	0.33	0.33	0.33	0.29	0.21	0.24	0.48	0.42	0.45	0.42	0.37	0.39	0.41	0.51	0.45	0.73	0.65	0.69
	Android	0.51	0.37	0.43	0.46	0.47	0.46	0.42	0.45	0.43	0.48	0.43	0.45	0.57	0.61	0.59	0.67	0.69	0.68
	Intel	0.35	0.35	0.35	0.36	0.32	0.34	0.47	0.45	0.46	0.46	0.43	0.45	0.58	0.57	0.57	0.71	0.71	0.71
	mixed	0.40	0.36	0.38	0.37	0.34	0.35	0.47	0.44	0.45	0.46	0.42	0.44	0.54	0.57	0.55	0.72	0.69	0.70

Table 3: Accuracy results in searching semantic similar functions

the victim AI program and a trained semantic representation model. We first parse the program to collect the graph files (CG) and collect all kernel nodes N in CG based on whether the node has an edge point. While list N is not empty, we select a node c from N based on whether c 's all precursors have been inferred (line 3). For the selected node c , we find the corresponding binary function f (line 4) and apply our semantic representation model to infer the kernel type k of this function (line 6). After that, we compute the output of node c and remove it from N . We repeat this process iteratively until N is empty. Finally, we return the computed tensor of the output node in CG . By doing so, we reconstruct the DNNs from the AI programs.

Algorithm 2 DNNs Reconstruction Algorithm

Input: VicProg. {Victim AI program}
Input: ϕ_{binary} {Semantic representation model}
1: $CG = \text{ParseFile}(\text{VicProg})$
2: $N = \text{CollectKernelNodes}(CG)$
3: **while** N is not empty **do**
4: $c = \text{SelectNode}(N)$ {select a node whose all precursors are computed}
5: $f = \text{FindBinaryFunction}(c)$
6: $k = \phi_{binary}(f)$ {Infer kernel type based on Eq. 4}
7: $c.type = k$ {Assign kernel to node c }
8: Compute node c
9: Remove c from N
10: **end while**
Output: Return computed tensor of the output node in CG .

Type	Platform	# Train	# Test	Instructions	Blocks
Vision	Android	22948	5737	194.75 ± 223.64	43.95 ± 47.19
	Intel	22892	5724	141.52 ± 135.41	42.49 ± 45.61
	ARM	28481	7121	208.19 ± 291.24	40.81 ± 44.70
	Aarch64	5761	1441	197.68 ± 216.46	43.87 ± 47.24
	Mix	80082	20023	195.61 ± 238.88	42.92 ± 46.41
Textual	Android	2647	662	131.48 ± 144.01	32.22 ± 38.23
	Intel	808	207	97.01 ± 104.40	32.96 ± 38.058
	ARM	2728	683	126.30 ± 142.88	21.20 ± 38.06
	Aarch64	2618	655	135.67 ± 150.20	32.30 ± 38.51
	Mix	8801	2207	127.94 ± 142.73	31.32 ± 37.66

Table 4: Data Statistics of Assembly Functions

4 Evaluation

We answer following questions through empirical evaluation.

- **RQ1:** How accurate of proposed representation model?
- **RQ2:** Can our algorithm 2 rebuild DNNs correctly?
- **RQ3:** Is our approach sensitive to hyper-parameters.
- **RQ4:** What's the contribution of each component?

4.1 Experimental Setup

Datasets. We first download the open source models from *TorchVision* and *HuggingFace* as our dataset. Our model dataset includes vision models (*e.g.* VGG-11, *etc.*) and textual models (*e.g.* BertSentenceClassification, *etc.*). We then use TVM, a popular deep learning compiler [Chen *et al.*, 2018b] to compile deep neural networks into AI programs. We consider four different deployment platforms (*e.g.*, Android, Intel, ARM, and Aarch64), and we set the compilation optimization level from 0 to 4. For the executable files of the compiled programs, we use radare2 to disassemble the binary functions into assembly instructions. We split 80% of

Model	# of Layers	# of Parameters	Acc Loss
LeNet-5	7	0.44M	0.00
ResNet-18	68	11.69M	0.00
ResNet-50	175	25.56M	0.00
MobileNet-V2	158	3.51M	0.00
VGG-11	29	132.86M	0.00
VGG-19	45	143.67M	0.00
WideResNet-50	175	68.88M	0.00

Table 5: Successfully reversed DNNs without accuracy loss

the data as our training dataset and the rest 20% for testing. The statistic of our dataset can be found in Table 4.

Comparison Baselines. We compare our proposed method with six unsupervised semantic representation methods: *word2vec*, *doc2vec*, and *asm2vec*. For *word2vec*, we use the average vectors of all instructions in the assembly function to represent the semantic. For *word2vec* and *doc2vec*, we consider both *Skip-Gram* and *CBoW* models. For a fair comparison, for each method, we set the dimensions of the embedding as 100.

4.2 (RQ1) Results for Binary Function Searching

We follow existing work [Alon *et al.*, 2018; Alon *et al.*, 2019] to apply *Semantic Precision*, *Semantic Recall* and *Semantic F-1* score as our metrics. Table 3 summarizes the main results. From the results, we could observe that: the proposed method significantly improves the binary function mapping accuracy across two datasets on all hardware platforms, The observation demonstrates our embedding technique’s success in enhancing the prediction capacity in the binary mapping problem.

4.3 (RQ2) Reconstruct DNNs

In this section, we evaluate the effectiveness of our reconstruction algorithm (*i.e.* Algorithm 2). Because a learning-based technique cannot guarantee soundness results, we choose DNNs with high kernel operator similarity scores as examples. The selected model architectures are listed in Table 5. We use the accuracy loss as our evaluation metric. Formally, we feed random inputs into the original and reconstructed models. Then, the accuracy loss is defined as the accuracy difference between the accuracy of the original model and the reconstructed model. Table 5 lists some examples of our successful reversed DNN models.

4.4 (RQ3) Parameter Sensitivity

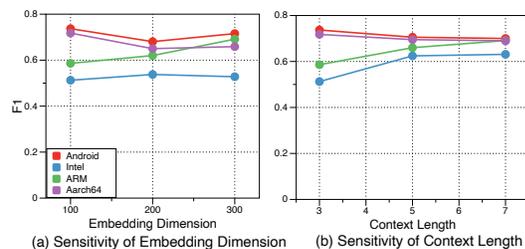


Figure 5: Sensitivity of Representation Model

Platform	Vision		Textual	
	Our	Baseline	Ours	Baseline
Android	0.21	0.80	0.09	0.71
Intel	0.11	0.60	0.62	0.85
ARM	0.12	0.65	0.51	0.98
Aarch64	0.14	0.75	0.08	0.66
Avg	0.14	0.70	0.32	0.80

Table 6: The OOV ratio results

We perform the sensitivity analysis of our proposed techniques on the Textual dataset. **(a) embedding dimension:** We first conduct experiments to understand how different embedding dimensions affect our technique. Specifically, we configure the embedding dimension as 100, 200, 300 and measure F-1 under different settings. **(b) context length:** We then configure the context size as 3, 5, 7 and train different models. The results of different embedding configuration are listed in Fig. 5. From the results, we observe that *NNReverse* can keep a stable F1 scores on different embedding dimension and context length settings, which indicates that *NNReverse* is not sensitive to embedding dimension and context length.

4.5 (RQ4) Ablation Studies

Can fine-grained modeling eliminate Out of Vocabulary (OOV)? We first conduct experiments to explore whether our fine-grained instruction embedding can eliminate the issue of OOV. We model each assembly instruction as a “word” and measure the OOV ratio, and then we measure the OOV ratio of our instruction embedding. The OOV results are summarized in Table 6. We observe that directly modeling each instruction as “word” will result in a high OOV ratio. This is because assembly instruction may contain the memory address, and the value of the address is had to predict.

The performance of each representation component? In this study, we compare our embedding method to solo text embedding and solo topological embedding to investigate the efficacy of embedding combinations. From the results in Table 2, we observe that: (1) for most cases, 20 out of 30, combining text embedding and structure embedding can increase binary mapping accuracy. (2) structure embedding performs the poorest in all scenarios, implying that simply applying structure information is insufficient to infer function semantics. CFG structures, on the other hand, are complementary to text, and combining them yields better results.

5 Discussion & Conclusion

In this research, we present a technique for automatically reversing DNNs from deployed AI applications. In particular, we propose a novel embedding technique for capturing the semantics of assembly functions. The experiments demonstrate that *NNReverse* is possible to correctly reverse DNNs.

Acknowledgments

This work was supported by NSF grant CCF-2146443.

References

- [Alon *et al.*, 2018] Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. code2seq: Generating sequences from structured representations of code. *arXiv preprint arXiv:1808.01400*, 2018.
- [Alon *et al.*, 2019] Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. code2vec: Learning distributed representations of code. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–29, 2019.
- [Carlini and Wagner, 2017] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [Carlini *et al.*, 2020] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. *arXiv preprint arXiv:2012.07805*, 2020.
- [Chen *et al.*, 2018a] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. {TVM}: An automated end-to-end optimizing compiler for deep learning. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 578–594, 2018.
- [Chen *et al.*, 2018b] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. TVM: An automated end-to-end optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 578–594, Carlsbad, CA, October 2018. USENIX Association.
- [Cyphers *et al.*, 2018] Scott Cyphers, Arjun K Bansal, Anahita Bhiwandiwalla, Jayaram Bobba, Matthew Brookhart, Avijit Chakraborty, Will Constable, Christian Convey, Leona Cook, Omar Kanawi, et al. Intel ngraph: An intermediate representation, compiler, and executor for deep learning. *arXiv preprint arXiv:1801.08058*, 2018.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [Duddu *et al.*, 2018] Vasisht Duddu, Debasis Samanta, D Vijay Rao, and Valentina E Balas. Stealing neural networks via timing side channels. *arXiv:1812.11720*, 2018.
- [Hua *et al.*, 2018] Weizhe Hua, Zhiru Zhang, and G Edward Suh. Reverse engineering convolutional neural networks through side-channel information leaks. In *2018 55th Design Automation Conference*, 2018.
- [Jagielski *et al.*, 2020] Matthew Jagielski, Nicholas Carlini, David Berthelot, Alex Kurakin, and Nicolas Papernot. High accuracy and high fidelity extraction of neural networks. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 1345–1362, 2020.
- [Krishna *et al.*, 2019] Kalpesh Krishna, Gaurav Singh Tomar, Ankur P Parikh, Nicolas Papernot, and Mohit Iyyer. Thieves on sesame street! model extraction of bert-based apis. *arXiv preprint arXiv:1910.12366*, 2019.
- [Li *et al.*, 2020] Mingzhen Li, Yi Liu, Xiaoyan Liu, Qingxiao Sun, Xin You, Hailong Yang, Zhongzhi Luan, Lin Gan, Guangwen Yang, and Depei Qian. The deep learning compiler: A comprehensive survey. *IEEE Transactions on Parallel and Distributed Systems*, 32(3):708–727, 2020.
- [Oh *et al.*, 2019] Seong Joon Oh, Bernt Schiele, and Mario Fritz. Towards reverse-engineering black-box neural networks. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pages 121–144. Springer, 2019.
- [Raffel *et al.*, 2019] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
- [Roberts *et al.*, 2019] Nicholas Roberts, Vinay Uday Prabhu, and Matthew McAteer. Model weight theft with just noise inputs: The curious case of the petulant attacker. *arXiv preprint arXiv:1912.08987*, 2019.
- [Rotem *et al.*, 2018] Nadav Rotem, Jordan Fix, Saleem Abdulrasool, Garret Catron, Summer Deng, Roman Dzhabarov, Nick Gibson, James Hegeman, Meghan Lele, Roman Levenstein, et al. Glow: Graph lowering compiler techniques for neural networks. *arXiv preprint arXiv:1805.00907*, 2018.
- [Shervashidze *et al.*, 2011] Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9), 2011.
- [Tramèr *et al.*, 2016] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, pages 601–618, 2016.
- [Yan *et al.*, 2020] Mengjia Yan, Christopher W Fletcher, and Josep Torrellas. Cache telepathy: Leveraging shared resource attacks to learn {DNN} architectures. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 2003–2020, 2020.
- [Zhu *et al.*, 2021] Yuankun Zhu, Yueqiang Cheng, Husheng Zhou, and Yantao Lu. Hermes attack: Steal {DNN} models with lossless inference accuracy. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021.