

On the Role of Pre-trained Embeddings in Binary Code Analysis

Alwin Maier
Max Planck Institute for
Solar System Research
Germany

Felix Weißberg
Technische Universität Berlin
Germany

Konrad Rieck
Technische Universität Berlin
& BIFOLD
Germany

ABSTRACT

Deep learning has enabled remarkable progress in binary code analysis. In particular, pre-trained embeddings of assembly code have become a gold standard for solving analysis tasks, such as measuring code similarity or recognizing functions. These embeddings are capable of learning a vector representation from unlabeled code. In contrast to natural language processing, however, label information is not scarce for many tasks in binary code analysis. For example, labeled training data for function boundaries, optimization levels, and argument types can be easily derived from debug information provided by a compiler. Consequently, the main motivation of embeddings does not transfer directly to binary code analysis.

In this paper, we explore the role of pre-trained embeddings from a critical perspective. To this end, we systematically evaluate recent embeddings for assembly code on five downstream tasks using a corpus of 1.2 million functions from the Debian distribution. We observe that several embeddings perform similarly when sufficient labeled data is available, and that differences reported in prior work are hardly noticeable. Surprisingly, we find that end-to-end learning *without* pre-training performs best on average, which calls into question the need for specialized embeddings. By varying the amount of labeled data, we eventually derive guidelines for when embeddings offer advantages and when end-to-end learning is preferable for binary code analysis.

ACM Reference Format:

Alwin Maier, Felix Weißberg, and Konrad Rieck. 2023. On the Role of Pre-trained Embeddings in Binary Code Analysis. In *Proceedings of 19th ACM ASIA Conference on Computer and Communications Security (ASIACCS'24)*. ACM, New York, NY, USA, 16 pages. <https://doi.org/https://doi.org/10.1145/3634737.3657029>

1 INTRODUCTION

Deep learning has been a driving force behind several advances in computer security. In particular, the ability of neural networks to distill information from highly complex data, such as assembly code, has led to a number of learning-based methods for binary code analysis. These methods allow, for example, to locate function boundaries [3, 30, 33], differentiate optimization levels [6, 31], assess

code similarity [26, 32, 37, 39], reconstruct arguments [7, 18, 19], and detect aliases in memory [17]. While the approaches differ in the architecture of the neural networks used, most share a key component: an *embedding*. This learned vector representation originates from the area of natural language processing and provides geometric access to the data's structure and semantics, forming a versatile basis for solving different learning tasks.

Over recent years, several methods have emerged for crafting embeddings tailored to binary code analysis, including *Asm2Vec* [11], *Instruction2Vec* [21], and *PalmTree* [23]. Additionally, dedicated approaches such as *Gemini* [36] and *SAFE* [25] have been specifically designed for generating function embeddings used to detect similar functions. The underlying rationale for these embeddings lies in their *pre-training* on large collections of unlabeled code, which allows for encoding general characteristics of the data and learning a versatile representations for various downstream tasks. Over time, these embeddings have become a gold standard for applying deep learning to binary code analysis [2, 4, 19, 29, 38, 39].

Although natural language processing bears striking similarities with code analysis, the availability of labeled data differs significantly between the two domains. For natural language text, tedious manual labeling is often unavoidable to create examples for supervised learning, rendering pre-trained embeddings indispensable. In contrast, for many tasks of binary code analysis, labeled data can be easily generated from debug information provided by a compiler. For example, labels for function boundaries, optimization levels, and argument types can be extracted during the compilation process and enable constructing large-scale training sets with label information automatically. As a result, the necessity of pre-trained embeddings in natural language processing does not naturally apply to tasks in binary code analysis, where end-to-end learning is frequently possible.

In this paper, we thus investigate the role of pre-trained embeddings for binary code from a critical perspective. For this investigation, we construct a labeled evaluation corpus of 1.2 million functions from the Debian distribution, totaling about 129 million x86 instructions. This corpus allows us to systematically evaluate the capabilities and limitations of five widely used embeddings for assembly code, namely *Word2Vec* [28], *Asm2Vec* [11], *Instruction2Vec* [21], and *PalmTree* [23]. In particular, we evaluate the performance of each embedding in different configurations on five common downstream tasks of binary code analysis: compiler detection, optimization level identification, function argument prediction, argument type reconstruction and code similarity detection.

Our results provide a new view on pre-trained embeddings in binary code analysis: First, we observe that the embeddings hardly differ in performance if sufficient training data is available for the

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

ASIACCS'24, July 1–5, 2024, Singapore

© 2023 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/https://doi.org/10.1145/3634737.3657029>

downstream task, so that differences discussed in prior work are not noticeable [23]. Even a random embedding of the instructions provides a reasonable performance in our experiments. Second, we surprisingly find that end-to-end learning *without* a pre-trained embedding yields the best performance on average. Contrary to our intuition, pre-training does not generally unlock additional information, and it provides no advantage in binary code analysis when sufficient labels are available.

Our work should not be interpreted as a general criticism of pre-trained embeddings. Instead, by reducing the amount of labeled data available, we can also demonstrate the utility of this technique in our evaluation. In particular, PalmTree [23] provides the best overall performance when labeled training data becomes scarce. We can thus derive guidelines to help practitioners decide whether or not to use pre-training. Consequently, our work adds a new facet to research on deep learning for binary code analysis, indicating that benefits from other domains do not necessarily carry over and need be critically reflected upon in practice. In summary, we thus make the following contributions:

- (1) *Critical evaluation of pre-trained embeddings.* We present a systematic evaluation of embeddings for binary code analysis with varying training data and embedding dimensions while also considering computational expenses.
- (2) *Large-scale evaluation corpus.* We provide researchers with an open corpus of 1.2 million labeled functions and 129 million x86 instructions from the Debian distribution for five downstream tasks (<https://github.com/a0x77n/orbit-dataset>).
- (3) *Recommendations for binary code analysis.* We derive guidelines to study the performance of pre-trained embeddings on assembly code and decide when to rely on conventional end-to-end learning instead.

Roadmap. We briefly review the background of pre-training and embeddings for assembly code in Section 2. Our benchmark corpus of labeled code for five downstream tasks is then introduced in Section 3 and the corresponding experiments in Section 4. We discuss our findings on pre-trained embeddings and derive recommendations in Section 5. Finally, Section 6 concludes the paper.

2 A PRIMER ON PRE-TRAINING

Let us start by introducing some background on the concept of pre-training and embeddings for assembly code, before critically reflecting on their role in binary code analysis.

2.1 Training with a Headstart

When humans learn a new task, such as a playing a music instrument, they usually do not start from scratch but are able to built upon prior knowledge. For example, someone who played the violin is probably faster in learning to play the cello than someone without prior experience. This insight into the human learning process has also been employed in the machine-learning domain and is commonly referred to as transfer learning.

To be more specific, a learning model trained for a certain task can help create a second model for a different task. The first task is commonly referred to as *pre-training task* and the second as *downstream task*. The rationale underlying this transfer is that the

second model may perform better when building on the knowledge of the first one. Since this transfer learning typically revolves around improving the performance of downstream task, we refer to the pre-training as task-agnostic (with respect to the downstream task), while the downstream training is task-specific.

Pre-training is considered especially beneficial in cases where high quality labelled data is expensive to create but unlabelled data is widely available. In natural language processing and computer vision, for example, vast amounts of unlabeled data are available on the Internet and can be employed for pre-training. Since the pre-training task and the downstream task do not need to share the same training objective, it becomes possible to use unsupervised learning for pre-training on a large unlabeled dataset and perform supervised learning on a smaller labeled one. A well-known example of this strategy is the unsupervised learning of input representations. In this case, the pre-training task learns a vector representation of the data, denoted as *embedding*, which serves as input for the subsequent downstream tasks, as shown in Figure 1(a).

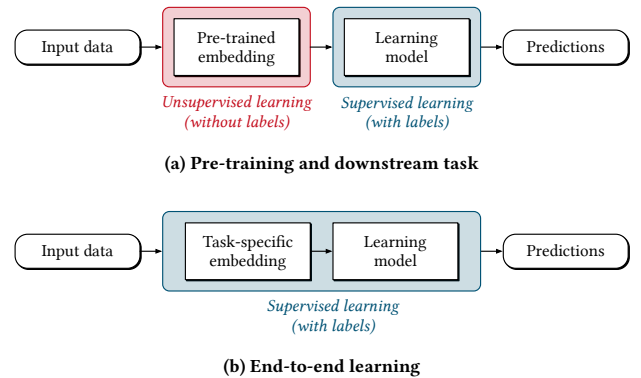


Figure 1: Schematic comparison of pre-training task with downstream task and conventional end-to-end learning.

This type of pre-training has been shown to be effective for several learning tasks in several domains. For example, in computer vision, pre-training has been employed for object recognition [12] and semantic segmentation [24]. Similarly, in natural language processing, embeddings have been successfully applied for question answering [10], machine translation [14], and text summarization [1]. Naturally, these advances have spawned a research in security aimed at improving binary code analysis. The rationale of this work has been to create *embeddings* for assembly code that provide a versatile representation, simplifying downstream tasks, such as function recognition and type inference.

Pre-training contrasts with so-called *end-to-end learning*. In this setup, the representation of the data is learned along with the task at hand, as shown in Figure 1(b). The previously separate tasks now focus on the same objective and work with the same labeled data. That is, the learned representation becomes task-specific and is only suitable for the particular downstream task. Unlabeled data cannot be used in this setting. Consequently, *end-to-end learning* requires access to a sufficient amount of labeled data.

2.2 Instruction Embeddings

We proceed to investigate the different embeddings that have been proposed for binary code analysis. To this end, we first give a formal definition and then provide a recap of recent approaches for embedding assembly code.

Simply put, embeddings are functions mapping data of the input domain to a vector representation. This vector can then be used for subsequent machine-learning tasks. The main reason behind using such a representation is, that it does not simply encode a word as a low-dimensional vector but is also able to express relationships among different words. For example, synonyms may be mapped to vectors within close proximity, while opposites are a mapped far apart from each other. Since these are appealing characteristics for machine learning, embeddings are widely used for natural language tasks [1, 10, 14] and recently for binary code analysis, where instructions take up the role of words.




Figure 2: Example of x86-64 assembly instructions.

Before introducing these embeddings, we first need to formally define what an instruction is. In particular, we refer to an instruction as the disassembled representation of a machine code operation for a given architecture. Generally, an instruction consists of an operation type denoted as *mnemonic* (M) and a group of *operands* (O_1, \dots, O_n). Typically, the operand of an instruction can have one of the following types:

- (1) *Register*. The operand is a processor *register*, where each register has a unique identifier.
- (2) *Immediate value*. The operand is an address value or a numerical constant.
- (3) *Memory access*. The operand is a memory access expression which usually consists of several components.

While the specific structure of the operands and the instruction depends on the architecture and the syntax flavor, the general structure remains the same. To give an example for a specific case, Figure 2 shows a sequence of five instructions for the x86-64 architecture in Intel syntax. For the instruction in line three, the mnemonic is highlighted in red and the operands in yellow. The first operand is a register operand, while the second one is a memory access expression. In this specific case, ①, ②, ③ and ④ point to the components of the memory access: the *base register*, *index register*, *scale* and *displacement value*, respectively. Although these instructions actually describe program semantics, it is easy to see that they can be interpreted as words (token) or phrases, similar to natural language text.

Embeddings for binary code differ in how they characterize the content of the instructions. While most embeddings build on tokenization schemes for operations and operands that are architecture independent, some approaches utilize specific knowledge

about instruction semantics and execution behavior. The embeddings considered in our work fall into the first category, with the exception of *Instruction2Vec*, which incorporates knowledge about the structure of x86-specific memory access expressions.

To describe common instruction embeddings jointly, we phrase the previous description more formally. To this end, let \mathcal{I} be the set of all instructions. The embedding method can then intuitively be described as a function Ψ mapping an instruction $I \in \mathcal{I}$ to a real valued vector of dimension d :

$$\Psi : \mathcal{I} \longrightarrow \mathbb{R}^d$$

Although sufficient for most cases, this definition is not able to capture methods which create context-dependent embeddings, such as *PalmTree*. That is, for a sequence of instructions $S = (I_1, \dots, I_m)$, the embedding for one instruction also depends on other instructions within the same sequence. To also capture such embedding methods we extend our definition to sequences of instructions. Let \mathcal{S} be the set of all instruction sequences $S = \bigcup_i I^i$. This allows for the following embedding function definition

$$\Phi : \mathcal{S} \longrightarrow \bigcup_i \mathbb{R}^{d \times i}$$

where a whole instruction sequence is mapped to a matrix in which each column represents the embedding vector of the respective instruction. With this definition it is possible to map the same instruction to different vector representations dependent on the surrounding instructions.

The remainder of this section introduces the six instruction embeddings used in our evaluation.

(a) *Word2Vec*. This embedding [25] is a variant of the classic embedding introduced by Mikolov et al. [28], which has been originally designed for natural language text. To learn a vector representation of assembly instructions, one can directly transfer this idea to binary code by considering instructions sequences as sentences and instructions themselves as words. The resulting embedding is based on the word embedding function ϕ of the original model and yields the function

$$\Psi(I) = \begin{cases} \phi(I), & \text{if } I \text{ is in the vocabulary} \\ 0^d & \text{otherwise.} \end{cases}$$

That means out-of-vocabulary (OOV) tokens are mapped to the zero vector. For the considered *Word2Vec* based embeddings approach, full instructions are considered as a single token and no further tokenization is conducted. However, immediate values greater than 5000 are normalized, i.e. replaced by the word IMM. In line with Massarelli et al. [25], our implementation uses the skip-gram algorithm [27] with negative sampling [28]. A similar instruction embedding based on a *Word2Vec* model has also been used by Chua et al. [8] for learning function type signatures.

(b) *Instruction2Vec*. The *Instruction2Vec* embedding [22] utilizes a classic model as well. As opposed to *Word2Vec*, it is trained over individual instruction components gathered from instruction sequences. Based on the component-based embedding ϕ that maps a component to a m dimensional space, an embedding for instructions is assembled. The embedding vector of an instruction I can be

divided into 9 slots of size m . Slot 1 contains the embedding vector of the operation (M), slot 2 up to slot 5 contains the embedding of the first operand (O_1), and slot 6 to slot 9 contains the embedding of the second operand (O_2). If an operand is absent, the respective slots are set to 0^m . The layout of the slots of the operands depends on the type of the respective operand:

- (1) O_i is a register x :
 $\text{slot}_{1+4*i+1} = \phi(x)$
- (2) O_i is an immediate x :
 $\text{slot}_{1+4*i+2} = (0, \dots, 0, x)$
- (3) O_i is a memory access (a, b, c, d) :
 $\text{slot}_{1+4*i+1} = \phi(a)$, where a is the base register,
 $\text{slot}_{1+4*i+2} = (0, \dots, 0, d)$, where d is the displacement value,
 $\text{slot}_{1+4*i+3} = \phi(b)$, where b is the index register,
 $\text{slot}_{1+4*i+4} = \phi(c)$, where c is the scale

The final embedding vector of I is obtained by concatenation of the nine slots: $\Psi(I) = \text{slot}_1 \parallel \dots \parallel \text{slot}_9$. Hence, the *Instruction2Vec* embedding has the dimension constraint $d = 9m$. As can already be seen by the way in which the embeddings are created, *Instruction2Vec* uses a finer granularity compared to the *Word2Vec* approach. That is, the instruction is broken down into mnemonics and operands, including registers, immediates, and memory access elements: For example, the third instruction of Figure 2 would be tokenized into ADD RAX RSP RCX 2 0x8.

(d) *Asm2Vec*. Ding et al. [11] introduces *Asm2Vec* [11], an embedding technique for binary functions based on the document embedding method PV-DM [20]. Similar to the *Instruction2Vec* approach, *Asm2Vec* utilizes an embedding function ϕ , mapping instruction components to an m -dimensional vector. To enhance the embedding, the instructions's syntactic structure is considered:

$$\Psi(I) = \phi(M) \parallel \frac{1}{N} \sum_{n=1}^N \phi(O_n)$$

Here, an instruction embedding is defined as the concatenation of the mnemonic and the averaged embeddings of the operands. The tokenization process in *Asm2Vec* is identical to *Instruction2Vec* and the embedding dimension d is constrained to even numbers: $d = 2m$. *Asm2Vec* provides embeddings for functions as well as instructions when intermediate steps are considered. We thus investigate both types in our evaluation.

(d) *PalmTree*. This embedding [23] is the most recent assembly language model build on top of BERT [10]. Since BERT is based on the transformer architecture [34] its models are capable of producing different embeddings for the same word dependent on the word's context. *PalmTree* can do the same for instructions. Hence, different from previous approaches, this embedding is context-dependent. Also different from the previous approaches, *PalmTree* interprets instructions as sentences and instruction tokens as words. Compared to the other approaches, *PalmTree* uses an even more fine grained tokenization which considers every syntactic element of the instruction. For example, the third instruction in Figure 2 is tokenized into ADD RAX , [RSP + (RCX * 2) + 0x8].

To train the language model, two tasks from the BERT approach are reused, namely, masked language modeling and an adapted version of next sentence prediction. Also a third task is taken into

account which builds on the clearly documented semantics of instructions (i. e. source and destination are known) and aims at predicting def-use relationships between two instructions. The *PalmTree* embedding function Φ_{PalmTree} maps a sequence of instructions to a real valued matrix by performing a mean pooling of the hidden states of the second last layer of the transformer encoder. This gives the instruction-sequence embedding function

$$\Phi(S) = \Phi_{\text{PalmTree}}(S)$$

and a context-dependent embedding representations for the i -th instruction of this sequence by $\Phi_{\text{PalmTree}}(S)_i$.

(e) *Random*. As a baseline, we also consider a random embedding. This embedding is conceptually trivial and just uses a random projection to map instructions to real-valued vectors. It can be expressed as $\Psi(I) = r_I$ with r_I being uniformly sampled from \mathbb{R}^d for each instruction. Note that this embedding can be considered a lower bound for the utility of embeddings, as it randomly represents instructions in a vector space without any information inferred from real-word binary code.

(f) *End-to-end*. The conceptual antagonist to these pre-trained embeddings is *end-to-end* learning. In this case, a so-called embedding layer is added to the neural network used as learning model. This layer executes the mapping Ψ from tokens to vectors before the data is forwarded to subsequent layers in the model. Technically, Ψ is defined by an embedding matrix which can be thought of as a lookup table. Different from pre-trained embeddings, the entries of the embedding matrix are trainable parameters of the neural network, realizing a task-specific representation.

For our implementation of an embedding layer in end-to-end learning, we process each function using *Asm2Vec*'s tokenization scheme to split instructions into tokens first. Based on these tokens we adapt a vocabulary with 2048 features consisting of the 2046 most frequent tokens and two special tokens, one used for masking and the other represents the out-of-vocabulary (OOV) token. The OOV token combines the less frequent tokens. Subsequently all tokens of a function are encoded using a unique number for each vocabulary token, resulting in a sequence of numbers which are streamed into the embedding layer. The embedding layer turns each number into a d -dimensional embedding vector.

2.3 Function embeddings

In contrast to embeddings at the instruction level, function embeddings aim to generate a comprehensive vector representation of an entire binary function. Before we present different embeddings for binary functions, we formally describe our notion of a binary function and define the embedding function.

A binary function is represented by its control-flow graph and comprising basic blocks, i.e., continuous sequences of disassembled instructions interconnected by directed edges that signify the flow of control. Formally, let $C = (\mathcal{V}, \mathcal{N} : \mathcal{V} \rightarrow \{x : x \subset \mathcal{V}\})$ be a control-flow graph with a set of basic blocks \mathcal{V} , and a function \mathcal{N} that maps each basic block to its successors. We can then define an embedding function using the set of all control flow graphs C as

$$\Omega : C \longrightarrow \mathbb{R}^d$$

that maps an arbitrary control-flow graph C of a binary function to a d -dimensional vector. This embedding function outputs a singular vector representation, effectively condensing the intricate structure of a binary function into a point within a continuous vector space, enabling streamlined comparison and analysis of functions in binary code.

In the following, we present four function embedding methods, two of which leverage instruction embeddings, resulting in a total of 15 different embeddings for our evaluation.

(a) *Asm2Vec*. As highlighted in the preceding section, *Asm2Vec* constitutes an embedding technique for binary functions, drawing inspiration from the PV-DM model rooted in *Word2Vec* principles. In PV-DM, the model learns to predict a word based on both context words and a paragraph vector. Throughout the training of *Asm2Vec*, the vector representations of instruction components and entire functions undergo multiple updates. Following training, the embedding of an unknown function C is derived by applying the model through gradient descent, allowing for the inference of vector representations $\Omega(C)$.

(b) *Gemini*. *Gemini*, introduced by Xu et al. [37], employs a neural network to detect similar binary functions, building upon the *Structure2Vec* algorithm [9]. In *Gemini*, basic block-level features are iteratively aggregated into function-level features based on the binary function’s control flow. While the original *Gemini* implementation relies on manually selected features [see 16, 36] for each basic block, we further incorporate pre-trained instruction embeddings as well as the *end-to-end* and *random* embedding types to generate diverse initial representations for basic blocks.

Given the instruction sequence of a basic block $v = (I_1, \dots, I_m) \in \mathcal{V}$ and the instruction-sequence embedding function Φ , the initial representation is computed as

$$x_v^\Phi = \frac{1}{m} \sum_{i=1}^m \Phi(v)_{*,i},$$

where $\Phi(v)_{*,i}$ is the i -th column and the embedding vector of I_i . In total, we compute six different vector representations $\Omega(C)$ for each binary function $C = (\mathcal{V}, \mathcal{N})$ using the same iterative approach of *Structure2Vec* architecture.

(c) *SAFE*. Another method for detecting similar functions is *SAFE* [25], which relies on a self-attentive neural network. Unlike *Asm2Vec* and *Gemini*, *SAFE* operates without the need for a control-flow graph. Instead, it treats functions as flat linear sequences of instructions, employing a recurrent neural network. The instruction sequence undergoes initial embedding using the *Word2Vec* technique before being fed into the neural network. Additionally, we incorporate the previously introduced instruction embeddings. In total, we generate five distinct embedding types, aligning with the network architecture proposed by Massarelli et al.. For a given function $C = (\mathcal{V}, \mathcal{N})$ and the instruction-sequence embedding function Φ , the embedding vector $\Omega(C)$ by running the embedded instruction sequence through the neural network.

(d) *Ada-002*. Finally, we consider a large language model for embedding. OpenAI’s *text-embedding-ada-002* represents their latest model, succeeding task-specific embeddings and offering a unified

representation¹. Designed for versatile applications, including text similarity, it features a maximum context length of 8192 tokens and an embedding dimension of 1536. While not exclusively tailored to binary code, it draws strength from extensive training on a diverse corpus, making it a state-of-the-art, all-purpose embedding.

3 AN EMBEDDING BENCHMARK

The cornerstone of our evaluation is a vast open corpus of labeled assembly code designed for various downstream tasks. Prior to delving into the assessment of the embeddings under consideration, we provide an introduction to these tasks (Section 3.1) and outline our pipeline for generating large-scale datasets automatically to train embeddings and learn models (Section 3.2).

3.1 Downstream Tasks

For our evaluation, we select five common downstream tasks by which we measure the performance of the different embeddings. Each task addresses a different challenge in binary analysis [6, 8, 11, 25, 31, 36]. Note that we are not proposing new solutions for these tasks but rather recreate existing experiments to investigate the role of pre-trained embeddings.

Table 1 offers a summary of the five downstream tasks, presenting the number of trainable parameters for the learning model and the parameters exclusive to the embedding layer in our end-to-end approach, along with their corresponding dimensions. Notably, Task T5, focused on detecting similar code, does not require an additional learning model as it relies solely on the embedding vectors. It’s worth highlighting the substantial variation in the number of model parameters across tasks, providing an opportunity for experimentation with learning tasks of varying complexities in the domain of binary code analysis.

Compiler and optimization options. The first two tasks deal with the detection of compilers (Task T1) and the identification of optimization options (Task T2). Both tasks have been previously investigated by Pizzolotto and Inoue [31] and Chen et al. [6]. Our implementation is based on the network introduced by Pizzolotto and Inoue. They propose a shallow network as learning model that consists of a single LSTM layer with a terminal dense output layer. The LSTM layer has an output dimension of 256 and performs the actual learning. The final layer uses the softmax activation function to project the prediction on one of the output classes. To handle the different classes of each task, we use 2 output nodes for the compiler detection (*GCC* and *CLang*) and 4 nodes for the identification of optimization options (00, 01, 02, 03). Different from the approach by Pizzolotto and Inoue [31] we use disassembled instructions as input as opposed to the raw bytes of machine code instructions.

Function type signatures. The other two downstream tasks deal with the recovery of function type signatures. We choose the implementation published by Chua et al. [8] and reuse their network architecture. According to their method, we create two tasks for this problem. The first task predicts the number of function parameters (Task T3) and the second task aims at determining the data type of the first parameter (Task T4), such as `int` or `char *`.

¹<https://openai.com/blog/new-and-improved-embedding-model>

Table 1: Number of classes and model size for each downstream task. The right column shows the number of additional weights for the embedding in the end-to-end learning setup.

Task	Downstream application	# Classes	Embedding dim.	# Total weights	# Embedding weights
T1	Compiler	2	128	656,898 (100 %)	262,144 (67 %)
T2	Optimization option	4	128	657,412 (100 %)	262,144 (40 %)
T3	Number of parameters	10	128	1,350,666 (100 %)	262,144 (19 %)
T4	Parameter type	7	128	1,349,895 (100 %)	262,144 (24 %)
T5	Function similarity	2	64	–	–

The learning model consists of three sequential GRU layers configured with dropouts to avoid overfitting. The final layer has 10 nodes for the prediction of the number of arguments and 7 nodes for the prediction of the argument types [see 8].

Function Similarity. In addition to assessing the efficacy of function embeddings, we explore two distinct approaches for learning embeddings for function similarity detection. The first approach is based on *Gemini* [37], which leverages a graph neural network. Another method for identifying similar binary functions is *SAFE* [25], which is founded on a self-attentive neural network architecture. We re-implement the original *Gemini* network and adjust parameters for *SAFE* to align with the complexity of *Gemini*.

While this task doesn't inherently necessitate a learning model, as embedding vectors can be directly compared using cosine similarity, we train different embedding models using *Gemini* and *SAFE*, integrating pre-trained instruction embeddings.

3.2 Mining Debian Packages

For the training of instruction embeddings with unlabeled data and the automatic construction of a labeled dataset for training the downstream tasks, we deploy two data processing pipelines. The first pipeline processes *Debian* binary packages and emits the pre-trained embeddings. The second pipeline uses *Debian* packages as well, but also requires an embedding model to generate the labeled dataset used to train and evaluate the downstream tasks. Both processes are depicted in Figure 3.

We manually build each package from source for the x86-64 architecture using eight compilation combinations. Each compilation combination uses a different compiler (*GCC* or *CLang*) and one of four optimization levels (O0 to O3). Note that building a single *Debian* source package can produce multiple binary packages for the same compilation combination.

In total, we extract 1,293,205 functions consisting of 129,487,277 instructions from 480 binary *Debian* packages. More details are provided in Table 2, where we list the number of binary packages (with and without variants due to compilation combinations), functions, and instructions for each source package. Since the dataset generation processes binary *Debian* packages our dataset can be expanded by adding more packages. It is also possible to create datasets for different architectures.

Training and test splits. We split the dataset in training and test data at the level of *Debian* source packages. Consequently, all functions from a single source package will be placed either in the training set or test set. We never spread functions from one source

Table 2: Our dataset is based on 480 binary *Debian* packages compiled from 8 source packages.

	# Packages	# Functions	# Instructions
Packages for training:			
binutils	208 / 26	1,044,289	110,948,051
coreutils	8 / 1	101,402	6,585,304
diffutils	8 / 1	5889	529,309
findutils	16 / 2	9096	767,734
Total	240 / 30	1,160,676	118,830,398
Packages for testing:			
inetutils	88 / 11	19046	1684013
sg3-utils	24 / 3	9412	1512412
usbutils	8 / 1	1020	121158
util-linux	120 / 15	103051	7339296
Total	240 / 30	132,529	10,656,879
Full dataset:			
Total	480 / 60	1,293,205	129,487,277

package over training and test data. We believe that this is an important detail as functions can be shared not only between the binaries of a packages, but also between binary packages of a single source package. While functions may also be shared across source packages, we argue that these cases are rare and do not necessarily lead to overfitting. Instead, they reflect scenarios in which source code is simply reused, for example, through the bad practice of copying and pasting code snippets. Our data split procedure is different from the approaches used by Chua et al. who divide each open source project into training and test data or Pizzolotto and Inoue who randomly partition all functions of the whole dataset into training and test data. Our procedure aims at reducing the risk that information from the training data leak to the test data.

Table 2 also shows which packages we use for training and which we use for testing. The processing of training as well as test packages are displayed in Figure 3.

Unsupervised data. For pre-training embeddings, we create a generic function corpus. To do so, we iterate over the binary *Debian* packages selected for training and hand them over to the corpus generator. The corpus generator creates a indefinitely repeating

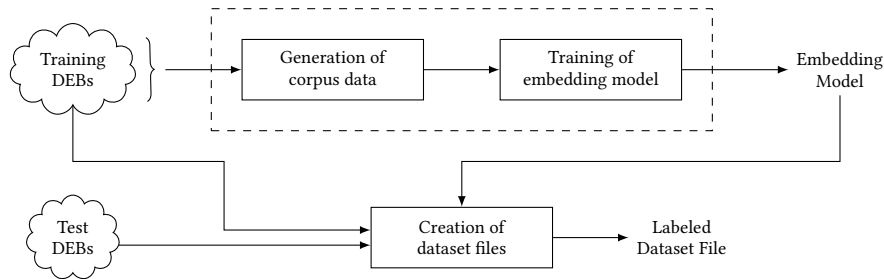


Figure 3: Dataset creation pipeline. A function corpus is build from all available training packages. Based on this corpus the embeddings are trained. Using the pre-trained embeddings, each *Debian* package is transformed into a dataset file containing the embedded functions and corresponding labels.

stream of disassembled and formatted instruction sequences on-the-fly, matching the requirements of the target embedding type. The streaming approach allows for processing large amounts of data without storing them in memory. Moreover, by working directly with *Debian* packages the corpus can be easily extended. This token stream is then used as unsupervised training data for the *Word2Vec*, *Instruction2Vec*, and *Asm2Vec* embeddings. It is important to note that packages selected for testing are not used for pre-training embeddings. This will prevent patterns from test packages to leak into the training process.

In particular, we extract all ELF binaries from each package. Next, we extract the opcodes of functions from the ELF binaries. Here, we leverage the DWARF [13] debugging information to identify every function and to pinpoint the address range of the functions to extract the corresponding opcodes from the text segment of the ELF binary. For this purpose, we use the python library *pyelftools*² to parse the text segment as well as the debugging information. Last, the opcodes are disassembled with the Capstone engine³ to produces a sequence of tokens for each function. Since each embedding type requires its specific instruction formatting and tokenization, we implement different tokenization schemes for the *Word2Vec*, *Instruction2Vec*, and *Asm2Vec* embeddings. Figure 3 shows where the corpus generator fits into the pipeline.

Embedding Pre-training. To pre-train the embedding models, we leverage the function corpus extracted from Debian packages, as illustrated in Figure 3. In this process, we make use of the *gensim*⁴ Python library, incorporating a module that implements the *Word2Vec* algorithm. The *Word2Vec* algorithm is applied to both the *Word2Vec* and *Instruction2Vec* embedding types, as both are grounded in a *Word2Vec* model, albeit with distinct encodings for each instruction. In the case of *Asm2Vec*, we extend *gensim* to support the *Asm2Vec* embedding, as detailed in Section 2.2 and Section 2.3. The training of instruction embeddings marks the concluding step in the pipeline depicted in Figure 3. Subsequently, these resulting embedding models are then reused to construct supervised datasets, facilitating the training and evaluation of various tasks, as well as the development of embedding models based on *Gemini* and *SAFE*.

Supervised data. To create a labeled dataset for training the downstream tasks and learning function embeddings we use the pre-trained embedding models to encode the instruction sequences of functions and the debug information to obtain the labels. We create an independent record file for each binary package and embedding. These dataset files contain the encoded functions and the corresponding labels for all four downstream tasks. Each file can then be loaded by TensorFlow⁵ and fits seamlessly into TensorFlow’s data processing pipeline.

To train embedding models based on the network architectures of *Gemini* and *SAFE*, labeled function pairs are required. While this presents a limitation for *Gemini* and *SAFE*, both approaches can incorporate pre-trained embeddings, allowing them to still benefit from unsupervised data in theory. All embedding models undergo training using a twin-network [5]. The dataset used for training the twin-networks is derived from the same Debian packages used to pre-train the instruction embeddings and *Asm2Vec*. This involves a specialized corpus generator sampling both positive function pairs (comprising the same functions with different compilation configurations) and negative function pairs (consisting of different functions from arbitrary packages).

All datasets used in our evaluation are available for other researchers: <https://github.com/a0x77n/orbit-dataset>

4 EXPERIMENTS

Equipped with a large corpus of labeled code for different downstream tasks, we are finally ready to compare the strengths and weaknesses of the considered embeddings. To this end, we design five basic experiments, each addressing a different aspect of pre-trained embeddings.

Experimental instances. For every experiment, we evaluate multiple *instances* of the downstream tasks. In this context, an instance is defined by the learning model of the task (T1-T4), the deployed embedding and the available supervised training data. For example, the learning model for task T1 paired with a *Word2Vec* embedding and all training data from Table 2 is one instance of the compiler identification task. The performance of each instance is then evaluated using the test data from Table 2. With this experimental setup, we can directly evaluate the capabilities of the embeddings, simply by comparing the instances considered in one experiment.

²<https://github.com/eliben/pyelftools>

³<https://www.capstone-engine.org/>

⁴<https://radimrehurek.com/gensim/>

⁵<https://www.tensorflow.org/>

Performance measures. As an evaluation criterion, we use the *accuracy* metric for tasks T1-T4, aligning with its prominent usage as the primary measure in the original publications [8, 31]. Task T5 employs the *area under the ROC curve*, a widely recognized metric for assessing binary classifiers, providing a comprehensive view of the trade-off between true-positive and false-positive rates. The *ROC curve* (receiver operating characteristic curve) [15] visually represents the performance of a classification model or detection system across various classification thresholds, plotting the true positive rate against the false positive rate.

Given a downstream task, the accuracy \mathcal{A}_F with respect to a set of labeled test functions F is given by the ratio of correct classified functions $f \in F$ to the total number of functions in F . The *general accuracy* \mathcal{A} of an instance is the performance over all test functions from Table 2. Analogously, we define the accuracy of a binary package. For this definition, we need a formal notion of a *compilation combination* c , which is an element of $C = \{GCC, CLang\} \times \{00, 01, 02, 03\}$. Now, given an instance of a downstream task, the *package accuracy* for the package p , denoted as \mathcal{A}_p , is defined as the average accuracy over all compilation combinations of p , that is, $\frac{1}{8} \sum_{c \in C} \mathcal{A}_F(p_c)$, where $F(p_c)$ denotes the functions in p build with compilation combination c . We make use of these notions of general and package accuracy for the different instances in our evaluation.

Experimental setup. Table 2 gives an overview of the datasets that are used for training and testing. Each instance is evaluated on 132,529 functions from 240 binary *Debian* packages each. The functions used for training are determined by the instance, consistently sourced from packages within the training set. In total, 1,160,676 functions from 240 packages are available for training. Notably, while we vary the amount of supervised training data in some experiments, the number of unsupervised training samples remains constant throughout. This means that pre-trained embeddings are consistently trained using the *entire* set of training data.

Whenever we consider an instance involving *end-to-end* learning, we add an additional dropout layer after the embedding layer to counteract overfitting. This layer randomly drops 20 % of its inputs. The reason for this extension is that unlike pre-trained embeddings, the size of the labeled data used to train instances with *end-to-end* learning varies, making them prone to overfitting when the labeled data is limited.

We proceed with the presentation of the five distinct experiments, each delineated in Sections 4.1 to 4.5. Within each section, the purpose of the experiment is elucidated alongside comprehensive details regarding its execution. It is worth noting that the reader has the flexibility to peruse these sections in any order they prefer. Furthermore, to facilitate navigation and comprehension, each section closes with a forward reference to the corresponding results expounded upon in Section 5. Additionally, for convenient reference, an overarching overview of all experiments is encapsulated in Table 3, providing a succinct summary for quick comprehension and navigation.

Table 3: Overview of the experiments. The right columns list the corresponding sections for the description of the experiments and results.

Experiment	# Shards	Dimension	Seq. len.	Exp.	Res.
Baseline	1	128 (126 ⁶)	512	4.1	5.1
Experiment 1	varies	128 (126 ⁵)	512	4.2	5.2
Experiment 2	varies	128 (126 ⁵)	512	4.3	5.3
Experiment 3	1	varies	512	4.4	5.4
Experiment 4	1	512	varies	4.5	A.3

4.1 Baseline Experiment

To begin, we introduce a general baseline experiment. That is, we design an experiment to test our implementations of the various embeddings and downstream tasks for correctness. We create five different instances for each combination of downstream task and embedding type and train them on all training data. Each instance is based on a different random seed in order to detect possible variations. We further use this experiment to revise the peculiarities of the downstream task and discuss the natural limits on the classification performance and identify variations in the test data. The results are presented in Section 5.1 and complemented in Appendix A.1.

4.2 Experiment 1: Size of Labeled Data

The objective of the first experiment is to investigate the effect of the amount of labeled data on the accuracy of the employed embeddings. Moreover, we consider the changes in training time. The rationale behind this experiment is the following: Unlike pre-trained embeddings, an end-to-end embedding is build from scratch and can only be trained with supervised data.

For each combination of downstream application and embedding type, we create nine different groups of instances. The groups differ in the amount of labeled data used for training. In contrast, the instances within one group are trained with the same number of functions, but the function sets are disjoint. On that account, we partition the training data T into n disjoint subsets $t_i^n, i = 1, \dots, n$, i.e., $T = \bigcup_{i=1}^n t_i^n$. We refer to each set t_i^n as a *shard* (of the training data). For instance, if $n = 1$ the same data that is used for pre-training the embeddings is used to train the instances. As n gets larger less supervised data is available for training the instances, e.g., for $n = 2^8$ less than 1 % of the data is used. Doing so, we mimic scenarios where more unlabeled than labeled data is available. We conduct this experiment for $n = 2^i, i = 0, \dots, 8$. In total, we train 511 ($\sum_{i=0}^8 2^i$) instances distributed over nine instance groups for all 24 combinations of downstream task and embedding type.

Note that the pre-trained embeddings remain unaffected by the value of n . They are consistently trained on the entire training set, while only the learning model of the downstream task is trained on reduced data, that is, it is trained on a single shard of the training data. In contrast, experimental instances that use *end-to-end* learning have only this single shard to infer all weights for the embedding layer as well as the learning model. Hence, the *end-to-end* approaches face a significant penalty for larger values of n . Our primary findings of this experiment are elaborated on in Section 5.2, while Appendix A.2 offers supplementary insights.

⁶The dimension of the *Instruction2Vec* embedding is 126.

4.3 Experiment 2: Computational Expense

In the second experiment, we investigate the computational overhead of *end-to-end* learning, assuming pre-trained models are readily available and thus only *end-to-end* learning incurs an additional run-time overhead. Using the setup from the previous section, we focus on the required epochs, that is, the number of training cycles needed for each instance.

Recall that we work with groups of instances that exist for each task and embedding pair, where group is defined by the size of the labeled data. For example, one group uses the entire training data while for another group it is split into multiple shards. In this case, the first group will contain a single instance and the other group contains instances for each shard. We compute the *mean number of epochs* for each group by averaging the number of required epochs of its instances. We present the results in Section 5.3.

4.4 Experiment 3: Embedding Dimension

The goal of the next experiment is to analyze the impact of the embedding dimension on the general accuracy of an instance. In case of an end-to-end approach the embedding dimension is an easy to change parameter of the network architecture. It can be tuned like any other hyper-parameter. By contrast, the dimension of a pre-trained embedding is fixed. If one wants to deploy a pre-trained embedding for some application, the user needs to pick the embedding dimension in advance. In other cases, a particular pre-trained embedding might only exist for a specific dimension. For that reason, it is worthwhile to determine if it is possible to identify a recommended value for a default embedding dimension that yields strong results across different downstream tasks.

We create several pre-trained embeddings with different embedding dimensions. We pick dimensions of the range from 1 to 256. Starting with the smallest possible dimension for each embedding we subsequently increase the dimension in each step. Some embeddings come with a restriction in regard to the dimension: The embedding dimension of *Instruction2Vec* is a multiple of 9 and the smallest possible dimension of *Asm2Vec* is 2. *PalmTree* has a fixed dimensionality of 128. Again, all instances are trained on the entire dataset. The results are summarized in Section 5.3.

4.5 Experiment 4: Sequence Length

In the last experiment, we seek to find a viable sequence length for the downstream tasks. We aim at determining a good compromise between training time and the general accuracy of an instance. This trade-off enables us to train and evaluate thousands of instances. To this end, we truncate each function sequence at various positions and observe the general accuracy of the instances. Starting with only the first instruction of each function, we double the sequence length subsequently. We stop at a maximum of 256 instructions per function. Figure 13 shows a bar plot of the distribution of function lengths. Most functions are rather short. Only few functions are longer than 256 instructions. These functions are accumulated in the last bar of the plots. We use the *PalmTree* embedding for this experiment, as it is the only context-dependent one. We report the result in Section A.3 of the appendix.

Table 4: General accuracy of the baseline experiment. The second value (\pm) reflects the minimal standard deviation observed for different random seeds, underscoring the stability of the results over all instances.

	Task T1	Task T2	Task T3	Task T4	Average
Instruction embeddings:					
<i>Word2Vec</i>	0.97 \pm .00	0.68 \pm .00	0.80 \pm .01	0.92 \pm .05	0.84 \pm .01
<i>Instr.2Vec</i>	0.95 \pm .00	0.67 \pm .01	0.81 \pm .01	0.92 \pm .09	0.84 \pm .01
<i>Asm2Vec</i>	0.97 \pm .00	0.68 \pm .01	0.83 \pm .00	0.93 \pm .02	0.85 \pm .00
<i>PalmTree</i>	0.97 \pm .00	0.69 \pm .01	0.86 \pm .00	0.94 \pm .07	0.86 \pm .01
<i>end-to-end</i>	0.96 \pm .00	0.68 \pm .01	0.88 \pm .01	0.94 \pm .05	0.87 \pm .01
<i>random</i>	0.97 \pm .00	0.67 \pm .00	0.88 \pm .01	0.94 \pm .05	0.87 \pm .00
Function embeddings:					
<i>Asm2Vec</i>	0.93 \pm .00	0.65 \pm .00	0.43 \pm .00	0.85 \pm .03	0.72 \pm .00
<i>ada-002</i>	0.87 \pm .00	0.66 \pm .00	0.60 \pm .00	0.90 \pm .07	0.76 \pm .00

5 RESULTS AND DISCUSSION

We will now delve into the outcomes of the five experiments outlined in Sections 5.1 to 5.4. The setup for each specific experiment can be found in Section 4, where their objectives and designs are described. Supplementary findings are also provided in the appendix for further reference.

5.1 Baseline Experiment

The outcome of the baseline experiment is summarized in Table 4 and part of Figure 5. Table 4 lists the general accuracy of experimental instances that are trained on the entire dataset, that is, the same amount of data is available to pre-training and *end-to-end* learning. Our results are in line with the original publications and show that our re-implementations are correct. [23, 31].

In Figure 5, it can be observed that when utilizing 100% of the training data, the area under the ROC curve ranges from 0.79 to 0.92 for *Gemini* and 0.79 to 0.90 for *SAFE*. It is essential to highlight that our dataset differs from the evaluation data previously employed for function similarity [11, 26, 37] rendering direct comparisons challenging. Despite these discrepancies, discernible patterns emerge: models using *Word2Vec* and *PalmTree* demonstrate optimal performance, while the *end-to-end* approach yields comparable results to *Asm2Vec*. Conversely, *ada-002* exhibits the least favorable performance in this experiment.

Apart from the reproduced performance, we observe *no* notable differences between the embeddings. This result is particularly striking for the *random* embedding that performs as well as the specialized instruction embeddings. We conclude that the embedding type is not as important when sufficient labeled data is available for the downstream task. Furthermore, we observe that the results are stable across five runs with different random seeds. This is quantified by the standard deviation provided in Table 4. The deviation is at most 0.01. This suggests that the used optimization algorithms are well calibrated and the networks capacities are adequate for the data basis. As a result, the training algorithm is able to find model weights that produce stable results for each downstream task.

5.2 Experiment 1: Size of Labeled Data

The preceding results suggest that the amount of labeled data available is a key factor in the success of *end-to-end* learning. Therefore, we turn to the first experiment in which this quantity is varied. The corresponding results for tasks T1-T4 are shown in Figure 4, and Figure 5 shows the results for task T5.

The plots show the performance of the embeddings trained with different amounts of labeled data. This amount is defined by the number of shards used to partition the training data. For example, if 8 shards have been generated from the training data, only $\frac{1}{8}$ or 12.5 % of the labeled data is available for supervised training of the learning models or, in the case of task T5, the embedding models.

Expectedly, the overall performance improves with fewer shards; that is, the more supervised training data is available, the better the prediction. This trend can be observed for all embedding types and downstream tasks. At a certain point, however, this improvement saturates, so that increasing the amount of labeled data does not improve the performance. We conclude that the size of our evaluation corpus is sufficient for the downstream tasks and further extensions of it would not change the results significantly.

When we investigate the left part of the plots, that is, instances with limited labeled data, we observe different yet subtle performance differences. In the case of tasks T1 and T2, the pre-trained embeddings *Word2Vec*, *Asm2Vec*, and *PalmTree* lead to better results when less labeled data is used. The *end-to-end* approach requires more labeled data to attain a comparable performance. This confirms the expectation one has for pre-trained embeddings: Since a vector representation for the input is already learned, the neural network has fewer degrees of freedom and thus can better generalize even if only a few labeled examples are given. This observation is also confirmed by the performance of the *random* embedding. Despite its simplicity, it also provides reasonable accuracy, suggesting that it is not the embedding that matters, but rather the reduced number of trainable weights of the networks.

Our findings for tasks T3 and T4 are the other way around. For both tasks, the pre-trained embeddings *Word2Vec*, *Asm2Vec*, and *PalmTree* show no benefit compared to the *end-to-end* approach. Even with less data available, the *end-to-end* approach gives better results. For both tasks, pre-training negatively affects the performance, and we cannot identify a general advantage over conventional learning. This is further contradicted by the results of the *random* embedding, which works just as well, even when trained on limited training data.

The function embeddings employed in task T5 exhibit a slightly different trend: when utilizing minimal data, no discernible differences can be observed. However, with an increased volume of available data, *Gemini* models with *Word2Vec* and *PalmTree*, as well as the manual selected features of the original *Gemini* implementation, perform better than the *end-to-end*-based models and *Asm2Vec*, which have a similar performance. Such differences cannot be identified for *SAFE* embedding models. In this context, the *end-to-end* embedding type and all pre-trained embeddings exhibit similar results, whereas the performance of the *random* embedding is comparatively weaker. Surprisingly, the performance of *Instruction2Vec* shifts from being subpar in the context of *Gemini* to displaying a remarkably competent performance with the *SAFE* model. In both

scenarios, the *ada-002* embedding fails to achieve competitive results. Notably, in this task, the *random* embedding finally reveals its limitations. In the case of *Gemini*, this embedding type fails to derive any benefit from a larger pool of supervised training data.

We must conclude that for the considered downstream tasks, pre-training is not necessarily beneficial. When we examine the average performance on the right-hand side of Figure 4, *end-to-end* learning and *PalmTree* provide the best results, regardless of the amount of labeled data available. These results suggest that when developing a learning-based method for binary code analysis, experimenting with *end-to-end* learning is clearly a sound and likely successful design decision.

Minimizing the labeled data. In addition, Table 5 shows the general performance averaged over all instances that are trained with 256 shards. In this case, each shard contains only approximately 4500 labeled training samples. Further, the table shows the margin to the baseline experiment with approximately 1,160,676 labeled samples. Considering task T2, the accuracy of the *end-to-end* approach drops by 0.08 while the pre-trained approaches drop by 0.04 (*Word2Vec* and *PalmTree*), 0.6 (*Instruction2Vec*), and 0.05 (*Asm2Vec*) only. Similar observations can be made for task T1. As already observed, the tasks T3 and T4 behave differently, i.e., all embedding types lose accuracy just about equally. The limitation of labeled data has the most impact on task T3 with a loss of up to 0.30. Task T4, however, is barely affected, which is attracting our attention. We show in the appendix that a different metric gives a different picture in this case.

Table 5: General accuracy of the considered embeddings with a minimum of labeled data. The second value (↓) displayed represents the margin to the accuracy on the full dataset.

	Task T1	Task T2	Task T3	Task T4	Average
Instruction embeddings:					
<i>Word2Vec</i>	0.94↓.02	0.64↓.04	0.55↓.25	0.87↓.05	0.75↓.09
<i>Instr.2Vec</i>	0.83↓.12	0.62↓.06	0.53↓.30	0.82↓.09	0.70↓.14
<i>Asm2Vec</i>	0.93↓.04	0.63↓.05	0.59↓.25	0.90↓.02	0.76↓.09
<i>PalmTree</i>	0.94↓.03	0.65↓.04	0.59↓.26	0.87↓.07	0.76↓.10
<i>end-to-end</i>	0.89↓.07	0.61↓.08	0.57↓.29	0.88↓.05	0.74↓.12
<i>random</i>	0.87↓.09	0.58↓.10	0.60↓.30	0.89↓.05	0.73↓.13
Function embeddings:					
<i>Asm2Vec</i>	0.86↓.07	0.62↓.03	0.39↓.04	0.82↓.03	0.67↓.04
<i>ada-002</i>	0.71↓.06	0.50↓.16	0.33↓.27	0.83↓.07	0.59↓.16

5.3 Experiment 2: Computational Expense

While we’ve noted advantages of *end-to-end* learning in binary code analysis, they do come with a prize: Figure 6 shows the training time for tasks T1-T4. It stands out that the *end-to-end* approach needs the most epochs for training. When training on more labeled data, we notice a significant drop in the number of epochs. Especially for tasks T3 and T4, where the *end-to-end* learner has a comparable training time once the full dataset is used. The epochs needed for instances with pre-trained embeddings remain relatively constant.




Figure 4: General accuracy of the considered embeddings in relation to the available labeled training data (shards)

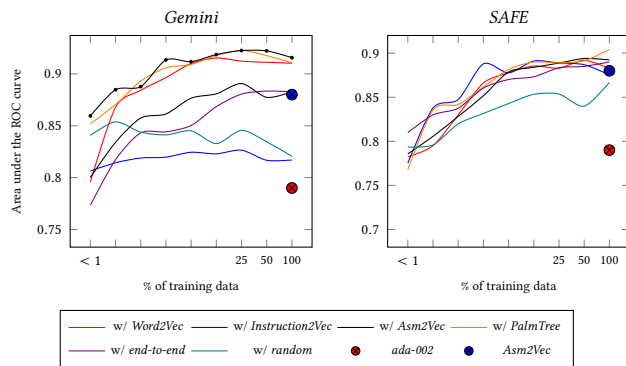


Figure 5: General accuracy of the considered embeddings in relation to the available labeled training data (shards).

We conclude that pre-trained embeddings have acquired relevant knowledge about the instructions, which helps training instances for each downstream application with fewer epochs.

An alternative perspective is presented for task T5 in Figure 7. Here, we maintain a consistent number of training samples per epoch, regardless of the available data. Consequently, training the embedding models for *Gemini* and *SAFE* with a larger volume of supervised data necessitates more epochs to cycle through all samples. Our observations indicate a general trend: the training algorithm effectively reduces the loss over more epochs when more data is available and concludes earlier with less data, yielding less favorable results, as depicted in Figure 5. Among the different embedding models, those based on *PalmTree* demonstrate the fastest training time while delivering one of the best results. Notably, the *Gemini* model with the *Instruction2Vec* embedding trains slowly, contrasting with its relatively short training time for *SAFE*. Unlike tasks T1-T4, the training time of *end-to-end*-based models is not as significantly impacted negatively; however, it is still noticeably slower.

We can conclude here that learning-based methods for code analysis can profit from a pre-trained embedding as the training process becomes faster. However, this is only relevant in scenarios where a complex network architecture is involved and the training time per epoch is high or the hardware resources are limited. Note that the prediction time is not affected, which is much more relevant in practice. We want to stress that we do not address the training

time of the pre-trained embeddings, since at least in theory those embeddings are only trained once, but used for many tasks.

5.4 Experiment 3: Embedding Dimension

It remains to investigate the role of the embedding dimension in our analysis, focusing specifically on tasks T1-T4 and the instruction embeddings. The results of this experiment are depicted in Figure 12 in the appendix. The plot shows the shift of the general accuracy of instances trained with different dimensions. Recall that the lowest dimension of *Asm2Vec* is 2 and *Instruction2Vec*'s embedding dimension is a multiple of 9. The *PalmTree* embedding, that we use as an off-the-shelf embedding, has a fixed dimension of 128.

This experiment shows that the *end-to-end* approach adapts better to lower embedding dimensions. On average, instances with an *end-to-end* embedding attain the best performance already with 16 dimensions. The pre-trained embedding types tend to require higher dimensionality. As a general rule, a dimension of 128 suffices for our tasks. This confirms the observation made by Li et al.

In summary, we have seen that an *end-to-end* embedding can get good results out of lower embedding dimensions. All pre-trained embedding types require a higher dimensionality to yield equal results. Our explanation is that pre-trained embeddings are more universal while an *end-to-end* embedding is always tailored towards the downstream task. Hence, an *end-to-end* embedding does not require the same complexity as pre-trained embeddings.

5.5 Discussion

Our experiments shed new light on the role of pre-trained embeddings in binary code analysis. We show that an *end-to-end* approach can compete with pre-trained embeddings if sufficient labeled data is available. If training time is not a constraint, *end-to-end* learning can obtain the same performance even with smaller embedding dimension. For pre-trained embeddings to be beneficial, the labeled data must be orders of magnitude smaller than the unlabeled data. This raises the question of whether these embeddings are actually relevant in practice, since the learning model of the downstream task requires labeled data in any case.

Moreover, our evaluation reveals that the embedding process is not crucial for the downstream task, as the success of the *random* embedding strikingly illustrates. We credit this finding to task-agnostic relations among instructions that are not as advanced as the connections between words in natural language. Without a given analysis task, an instruction embedding cannot carry the

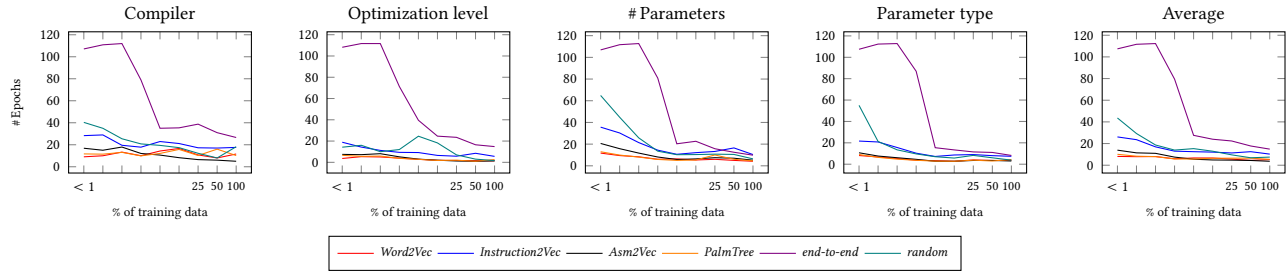


Figure 6: Training time of the considered embeddings in relation to the supervised data size.

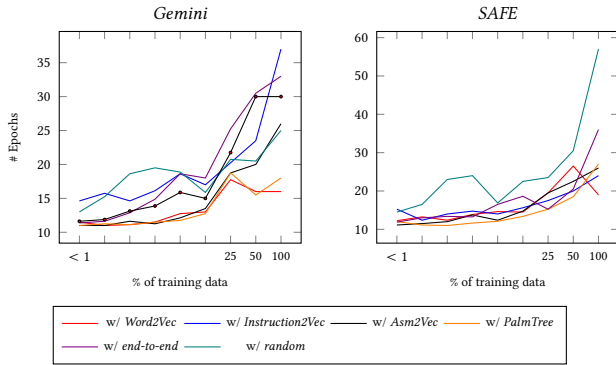


Figure 7: Number of epochs needed to train embeddings for code similarity in relation to the supervised data size.

same amount of information as a natural word embedding. This makes instruction embeddings less useful and also explains why the differences between our pre-trained embeddings are negligible.

Overall, our empirical analysis demonstrates that a general benefit of pre-training does not exist for binary code analysis in practice. This is contradictory to previous research in our domain [e.g., 23], and the common strategy to favor pre-trained embeddings over *end-to-end* learning. Interestingly, research in other fields also arrived at this observation [35] and questioned the role of pre-training if sufficient labeled information is available. Thus, our analysis refutes the intuition that pre-training is a generally beneficial and therefore mandatory step in designing methods for binary analysis.

Based on our observations, we work out the following recommendations for deep learning in binary code analysis.

- R1 Our first recommendation is to always try *end-to-end* learning, as it produces tailored embeddings for the task at hand and is easy to implement. Consequently, *end-to-end* learning should always be considered as the baseline approach.
- R2 When there is little labeled data or the size of the unlabeled data is orders of magnitude larger, the use of a pre-trained embedding is indicated. We recommend the use of the *PalmTree* model as it performs best in our evaluation.
- R3 We also recommend to differentiate between training and inference time. Only if the training time for a model is constrained, pre-trained embeddings are clearly preferable over slow *end-to-end* learning.

R4 Our last recommendation is to always publish pre-trained embeddings. If researchers need to re-train them anyway, their benefit vanishes and again *end-to-end* learning becomes a viable alternative.

Our recommendations are to be understood as guidelines and must be adapted to the respective downstream task. Nevertheless, we argue that pre-training is not a “silver bullet” in our field and conventional learning concepts must not be neglected from the start when developing new approaches for binary code analysis.

6 CONCLUSION

This paper investigates the role of pre-training on binary code analysis. To this end, we compare four pre-trained embeddings, an *end-to-end* approach, and a random instruction embedding under different downstream tasks. Our results show that binary analysis tasks with sufficiently labeled data do not benefit from pre-trained embeddings. Instead, conventional *end-to-end* learning provides the best performance on average. Only if the labeled data is artificially reduced, we can observe an advantage of pre-training.

Our results have consequences for applying deep learning in other tasks of binary code analysis. First, an *end-to-end* setup is typically easy to deploy along with the learning model. Hence, if labeled data is available, it should be the first option when developing a new approach. Second, when labeled data is scarce or the training time is constrained, pre-trained embeddings are a reasonable solution. For example, we could show that *PalmTree* often provides good results when we capped the available labeled data.

Overall, our study highlights an interesting aspect of interdisciplinary research. The benefits of embeddings in one research area can only be transferred to another if the experimental setup is limited by similar constraints. Once these constraints are lifted from a setup, the performance improvement may disappear. Therefore, we recommend that practitioners develop learning-based approaches judiciously and, when in doubt, prefer well-known learning concepts over the latest inventions.

ACKNOWLEDGEMENTS

This work was funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2092 CASA – (390781972), the European Research Council (ERC) under the consolidator grant MALFOY (101043410), and the German Federal Ministry of Education and Research under the grant BIFOLD24B.

REFERENCES

- [1] Armen Aghajanyan, Anchit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. 2021. Muppet: Massive Multi-task Representations with Pre-Finetuning. In *Proc. of the Conference on Empirical Methods in Natural Language Processing*. 5799–5811.
- [2] Sunwoo Ahn, Seongwan Ahn, Hyungjoon Koo, and Yunheung Paek. 2022. Practical Binary Code Similarity Detection with BERT-based Transferable Similarity Learning. In *Proc. of the Annual Computer Security Applications Conference (ACSAC)*. 361–374.
- [3] Jim Alves-Foss and Jia Song. 2019. Function boundary detection in stripped binaries. In *Proc. of the Annual Computer Security Applications Conference (ACSAC)*. 84–96.
- [4] Sajib Biswas, Timothy Barao, John Lazzari, Jeret McCoy, Xiuwen Liu, and Alexander Kostandarithes. 2022. Geometric Analysis and Metric Learning of Instruction Embeddings. In *Proc. of the International Joint Conference on Neural Networks (IJCNN)*.
- [5] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. Signature Verification using a “Siamese” Time Delay Neural Network. In *Advances in Neural Information Processing Systems*, Vol. 6.
- [6] Yu Chen, Zhiqiang Shi, Hong Li, Weiwei Zhao, Yiliang Liu, and Yuansong Qiao. 2018. HIMALIA: Recovering Compiler Optimization Levels from Binaries by Deep Learning. In *Proc. of the Intelligent Systems Conference (IntelliSys)*.
- [7] Zheng Leong Chua, Shiqi Shen, Prateek Saxena, and Zhenkai Liang. 2017. Neural Nets Can Learn Function Type Signatures From Binaries. In *Proc. of the USENIX Security Symposium*. 99–116.
- [8] Zheng Leong Chua, Shiqi Shen, Prateek Saxena, and Zhenkai Liang. 2017. Neural Nets Can Learn Function Type Signatures From Binaries. In *Proc. of the USENIX Security Symposium*. 99–116.
- [9] Hanjun Dai, Bo Dai, and Le Song. 2016. Discriminative Embeddings of Latent Variable Models for Structured Data. In *Proc. of the International Conference on Machine Learning (ICML)*. 2702–2711.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proc. of the Conference of The North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- [11] Steven Ding, Benjamin Fung, and Philippe Charland. 2019. Asm2Vec: Boosting Static Representation Robustness for Binary Clone Search against Code Obfuscation and Compiler Optimization. In *Proc. of the IEEE Symposium on Security and Privacy*.
- [12] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. 2014. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. In *Proc. of the 31st International Conference on Machine Learning*. 647–655.
- [13] DWARF Debugging Information Format Committee. 2010. *DWARF debugging information format*. DWARF Debugging Information Format Committee. Version 4.
- [14] Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding Back-Translation at Scale. In *Proc. of the Conference on Empirical Methods in Natural Language Processing*. 489–50.
- [15] Tom Fawcett. 2006. An introduction to ROC analysis. *Pattern recognition letters* 27, 8 (2006), 861–874.
- [16] Qian Feng, Rundong Zhou, Chengcheng Xu, Yao Cheng, Brian Testa, and Heng Yin. 2016. Scalable Graph-based Bug Search for Firmware Images. In *Proc. of the ACM Conference on Computer and Communications Security (CCS)*. 480–491.
- [17] Wenbo Guo, Dongliang Mu, Xinyu Xing, Min Du, and Dawn Song. 2019. DEEP-VSA: Facilitating Value-set Analysis with Deep Learning for Postmortem Program Analysis. In *Proc. of the USENIX Security Symposium*.
- [18] Jingxuan He, Pesho Ivanov, Petar Tsankov, Veselin Raychev, and Martin T. Vechev. 2018. Debin: Predicting Debug Information in Stripped Binaries. In *Proc. of the ACM Conference on Computer and Communications Security (CCS)*. 1667–1680.
- [19] Xin Jin, Kexin Pei, Jun Yeon Won, and Zhiqiang Lin. 2022. SymLM: Predicting Function Names in Stripped Binaries via Context-Sensitive Execution-Aware Code Embeddings. In *Proc. of the ACM Conference on Computer and Communications Security (CCS)*. 1631–1645.
- [20] Quoc Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *Proc. of the International Conference on Machine Learning (ICML)*.
- [21] 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 (2019).
- [22] 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 *Proc. of the International Conference on Internet (ICINI)*.
- [23] Xuezixiang Li, Yu Qu, and Heng Yin. 2021. PalmTree: Learning an Assembly Language Model for Instruction Embedding. In *Proc. of the ACM Conference on Computer and Communications Security (CCS)*.
- [24] Di Lin, Guangyong Chen, Daniel Cohen-Or, Pheng-Ann Heng, and Hui Huang. 2017. Cascaded feature network for semantic segmentation of RGB-D images. In *Proc. of the IEEE international conference on computer vision*. 1311–1319.
- [25] Luca Massarelli, Giuseppe Antonio Di Luna, Fabio Petroni, Roberto Baldoni, and Leonardo Querzoni. 2019. SAFE: Self-Attentive Function Embeddings for Binary Similarity. In *Proc. of the Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA)*.
- [26] Luca Massarelli, Giuseppe Antonio Di Luna, Fabio Petroni, Roberto Baldoni, and Leonardo Querzoni. 2019. SAFE: Self-Attentive Function Embeddings for Binary Similarity. In *Proc. of the Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA)*.
- [27] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *Proc. of the International Conference on Learning Representations (ICLR Workshop)*.
- [28] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems* 26 (2013).
- [29] Kexin Pei, Jonas Guan, Matthew Broughton, Zhongtian Chen, Songchen Yao, David Williams-King, Vikas Ummadisetty, Junfeng Yang, Baishakhi Ray, and Suman Jana. 2021. StateFormer: Fine-Grained Type Recovery from Binaries using Generative State Modeling. In *Proc. of the ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.
- [30] Kexin Pei, Jonas Guan, David Williams-King, Junfeng Yang, and Suman Jana. 2021. XDA: Accurate, Robust Disassembly with Transfer Learning. In *Proc. of the Network and Distributed System Security Symposium (NDSS)*.
- [31] Davide Pizzolotto and Katsuro Inoue. 2021. Identifying Compiler and Optimization Level in Binary Code From Multiple Architectures. *IEEE Access* 9 (2021).
- [32] Kimberly Redmond, Lannan Luo, and Qiang Zeng. 2019. A cross-architecture instruction embedding model for natural language processing-inspired binary code analysis. In *Proc. of the Workshop on Binary Analysis Research (BAR)*.
- [33] Eui Chul Richard Shin, Dawn Song, and Reza Moazzezi. 2015. Recognizing Functions in Binaries with Neural Networks. In *Proc. of the USENIX Security Symposium*. 611–626.
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*.
- [35] Sinong Wang, Madian Khabsa, and Hao Ma. 2020. To Pretrain or Not to Pretrain: Examining the Benefits of Pretraining on Resource Rich Tasks. In *Proc. of Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [36] 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 *Proc. of the ACM Conference on Computer and Communications Security (CCS)*.
- [37] 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 *Proc. of the ACM Conference on Computer and Communications Security (CCS)*. 363–376.
- [38] 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 *Proc. of the AAAI Conference on Artificial Intelligence*.
- [39] Fei Zuo, Xiaopeng Li, Patrick Young, Lannan Luo, Qiang Zeng, and Zhixin Zhang. 2019. Neural Machine Translation Inspired Binary Code Similarity Comparison beyond Function Pairs. In *Proc. of the Network and Distributed System Security Symposium (NDSS)*.

A SUPPLEMENTARY RESULTS

A.1 Baseline Experiment

	00	01	02	03
00	100.00 %	0.00 %	0.00 %	0.00 %
01	0.00 %	100.00 %	18.78 %	17.74 %
02	0.00 %	19.23 %	100.00 %	69.31 %
03	0.00 %	18.41 %	71.15 %	100.00 %

Figure 8: Overlap of functions compiled with different optimization levels. The arrows indicate the reading direction.

A closer look at optimization levels. The results for task T2 are worse compared to the other applications. We find that this decrease is not due to a more complex downstream task, but due to an

inherent problem in determining optimization levels. To understand this problem, we first recall that the optimization O3 activates the same flags as O2 and a few others. This means that the flags enabled by O2 are a subset of the flags enabled by O3. Unfortunately, not every function provides sufficient complexity to benefit from the unique flags of O3. As a result, both options may produce the same assembly instructions for some functions, making it impossible to recover to the original optimization level. The same is true for the other optimization level. However, the option O2 adds many more flags in addition to those added by O1 making the differentiation possible in most cases.

Figure 8 shows the overlap between the different optimization levels in our evaluation corpus. About 70 % of the functions compiled with option O2 could be compiled with option O3 just as well and vice versa. Less prevalent are functions compiled with option O1 that could also be compiled with options O2 and O3. Interestingly, the same issue also exists for the compiler-detection task. However, in this case it is far less pronounced and the overlap of the labels is not that common as can be seen in Figure 9. As a consequence, performance values as obtained for the other downstream tasks are impossible to achieve in this setting. Considering this observation, all embedding types actually produce results close to the optimum for the tasks T1 and T2.




Figure 9: Overlap of functions produced with different compilers. The arrows indicate the reading direction.

Analyzing the package performance. In addition to the previous observations, we want to check whether there are significant variations in prediction performance for different binary packages. For that purpose, we first compute the package accuracy of each instance. Second, we compute statistical characteristics for each downstream task and embedding type.

We depict the results as box-and-whiskers plots in Figure 10. Since we used five different seeds and our dataset contains 30 different binary packages for testing (without taking account of different compilation combinations), each plot is based on 150 values. The boxes cover the interquartile range (IQR) from the first to the third quartile. That means each box encompasses the mean accuracy values of 50 % of the binary packages. The plot shows that the prediction accuracy on 50 % of the binary packages differs by less than 5 % for tasks T1 and T2 and less than 10 % for tasks T3 and T4. The whiskers extend no more than 1.5 · IQR.

Overall, we observe that the majority of the packages fall within a normal range of packet accuracy. Nevertheless, there are some outliers that allow for better or worse performance on some of the downstream tasks. We investigate these outliers by identifying the individual Debian packages. We find that these packages are outliers for all considered embedding, indicating again that the choice of an embedding is less relevant to the underlying learning task when sufficient labeled data is available.

Table 6: General balanced accuracy of the considered embeddings with minimum of labeled data. The second value (↓) displayed represents the margin to the balanced accuracy achieved on the full dataset.

	Task T1	Task T2	Task T3	Task T4	Average
Instruction embeddings:					
<i>Word2Vec</i>	0.94 _{↓.02}	0.59 _{↓.04}	0.28 _{↓.32}	0.26 _{↓.28}	0.52 _{↓.17}
<i>Instr.2Vec</i>	0.83 _{↓.12}	0.57 _{↓.06}	0.30 _{↓.33}	0.23 _{↓.25}	0.48 _{↓.19}
<i>Asm2Vec</i>	0.93 _{↓.04}	0.59 _{↓.05}	0.31 _{↓.33}	0.31 _{↓.32}	0.53 _{↓.19}
<i>PalmTree</i>	0.94 _{↓.03}	0.61 _{↓.04}	0.31 _{↓.32}	0.29 _{↓.31}	0.54 _{↓.18}
<i>end-to-end</i>	0.89 _{↓.07}	0.56 _{↓.09}	0.27 _{↓.39}	0.29 _{↓.30}	0.50 _{↓.21}
<i>random</i>	0.87 _{↓.09}	0.53 _{↓.10}	0.30 _{↓.41}	0.30 _{↓.19}	0.50 _{↓.20}
Function embeddings:					
<i>Asm2Vec</i>	0.86 _{↓.07}	0.57 _{↓.04}	0.19 _{↓.06}	0.27 _{↓.20}	0.47 _{↓.09}
<i>ada-002</i>	0.70 _{↓.17}	0.44 _{↓.18}	0.11 _{↓.27}	0.20 _{↓.42}	0.36 _{↓.26}

A.2 Experiment 1: Size of Labeled Data

To check for variations in the general performance of instances trained on shards, we again use box-and-whiskers plots. Each plot is based on 256 different instances, each trained on a different data shard. We show the results of this experiment in Figure 11. While some variations exist, most instances produce similar results. We reason that the functions from some shards are not representative for the test data and no viable instance can be trained from those functions alone.

A.3 Experiment 4: Sequence length

Figure 14 shows the results of this experiment. It shows the general accuracy of instances trained with the *PalmTree* instruction embedding model in relation to the sequence length. That is, the maximum number of instructions of each function used to train the instances. The plot clearly shows that a certain number of instructions is needed. This comes at no surprise, however, the plot also shows that after processing about 16 instructions the performance starts to plateau and only little gain can be observed. We observe no further improvement after 128 instructions. Further, Figure 13 shows that only few functions are longer than 256 instructions. Because of that, we limit the instruction sequences in the other experiments to this threshold.

However, that such short sequences are sufficient raises questions, at least for tasks T3 and T4, since we expect that necessary information is missing in such short sequences. For that reason we measure the general accuracy with a different metric, the balanced accuracy, which takes the label imbalance into account. The results are also shown in Figure 14 and gives a different picture. The reason is, that now the underrepresented type categories have more weight. The classes of tasks T1 and T2 are more balanced and, hence, the performance is not affected.

Table 6 serves the same purpose as Table 5, but shows the performance based on the balanced accuracy. The average performance is noticeable lower. Moreover, the performance drop, that is, the margin to the performance obtained on the full dataset, is significant.




Figure 10: Package accuracy of instances trained on the full dataset presented as boxplots for all considered instruction embeddings across all binary packages. Outliers are visually highlighted as horizontal lines for easy identification.




Figure 11: Package accuracy of instances trained with a minimum of labeled data presented as boxplots for all considered instruction embeddings across all binary packages. Outliers are visually highlighted as horizontal lines for easy identification.




Figure 12: General accuracy of the considered embeddings in relation to the embedding dimension.

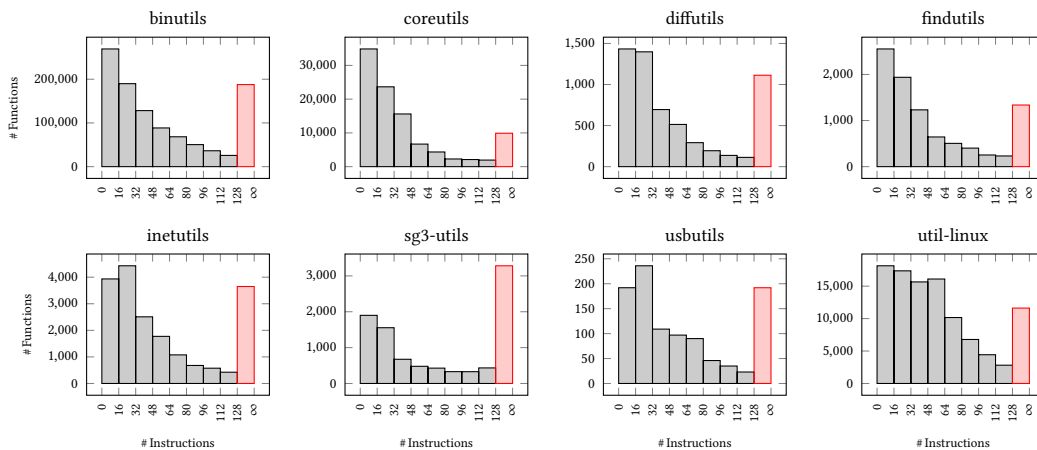


Figure 13: Distribution of function lengths. The red bar accumulates all functions with more than 128 instructions.




Figure 14: General accuracy for the *PalmTree* embedding in relation to the sequence length.