



The EarlyBIRD Catches the Bug: On Exploiting Early Layers of Encoder Models for More Efficient Code Classification

Anastasiia Grishina

anastasiia@simula.no
Simula Research Laboratory
Oslo, Norway

Max Hort

maxh@simula.no
Simula Research Laboratory
Oslo, Norway

Leon Moonen

leon.moonen@computer.org
Simula Research Laboratory &
BI Norwegian Business School
Oslo, Norway

ABSTRACT

The use of modern Natural Language Processing (NLP) techniques has shown to be beneficial for software engineering tasks, such as vulnerability detection and type inference. However, training deep NLP models requires significant computational resources. This paper explores techniques that aim at achieving the best usage of resources and available information in these models.

We propose a generic approach, EarlyBIRD, to build composite representations of code from the early layers of a pre-trained transformer model. We empirically investigate the viability of this approach on the CodeBERT model by comparing the performance of 12 strategies for creating composite representations with the standard practice of only using the last encoder layer.

Our evaluation on four datasets shows that several early layer combinations yield better performance on defect detection, and some combinations improve multi-class classification. More specifically, we obtain a +2 average improvement of detection accuracy on Devign with only 3 out of 12 layers of CodeBERT and a 3.3x speed-up of fine-tuning. These findings show that early layers can be used to obtain better results using the same resources, as well as to reduce resource usage during fine-tuning and inference.

CCS CONCEPTS

• **Software and its engineering**; • **Computing methodologies**
→ **Neural networks**; **Natural language processing**;

KEYWORDS

sustainability, model optimization, transformer, code classification, vulnerability detection, AI4Code, AI4SE, ML4SE

ACM Reference Format:

Anastasiia Grishina, Max Hort, and Leon Moonen. 2023. The EarlyBIRD Catches the Bug: On Exploiting Early Layers of Encoder Models for More Efficient Code Classification. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '23)*, December 3–9, 2023, San Francisco, CA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3611643.3616304>



This work is licensed under a Creative Commons Attribution 4.0 International License.

ESEC/FSE '23, December 3–9, 2023, San Francisco, CA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0327-0/23/12.

<https://doi.org/10.1145/3611643.3616304>

1 INTRODUCTION

Automation of software engineering (SE) tasks supports developers in creation and maintenance of source code. Recently, deep learning (DL) models have been trained on large open-source code corpora and used to perform code analysis tasks [3, 8, 27, 38]. Motivated by the naturalness hypothesis stating that code and natural language share statistical similarities, researchers and tool vendors have started training deep NLP models on code and fine-tuning them on SE tasks [11]. Amongst others, such models have been applied to type inference [17], code clone detection [50], program repair [9, 15, 47, 48], and defect prediction [7, 30, 35, 44]. In NLP-based approaches, SE tasks are frequently translated to code classification problems. For example, detection of software vulnerabilities is a binary classification problem, bug type inference is a multi-class classification setting, and type inference is a multi-label multi-class classification task in case a type is predicted for each variable in the program.

Most modern NLP models build on the transformer architecture [42]. This architecture uses *attention* mechanism and consists of an *encoder* that converts an input sequence to a *representation* through a series of layers, followed by *decoder* layers that convert this representation to an output sequence. Although effective in terms of learning capabilities, the transformer design results in multi-layer models that need large amounts of data for training from scratch. A well-known disadvantage of these models is the high resource usage that is required for training due to both model and data sizes. While a number of pre-trained models have been published recently, fine-tuning these models for specific tasks still requires additional computational resources [27].

This paper explores techniques that aim at optimizing the use of resources and information available in models during fine-tuning. In particular, we consider open white-box models, for which the weights from each layer can be extracted. We focus on encoder-only models, as they are commonly used for SE classification tasks, in particular, the transformer-based encoders. The standard practice in encoder models is to obtain the representation of the input sequence from the last layer of the model [14], while information from earlier layers is usually discarded [21]. I.e., while the early layers are used to *compute* the values of the last layer, they are generally not considered as individual representations of the input in the way that the last layer is. To exemplify the amount of discarded information at inference, when fine-tuning a 12-layered encoder, such as CodeBERT [14], for bug detection, 92% of the code embeddings are ignored.¹ However, it has been shown for natural

¹ That is, the weights from 11 out of 12 layers are ignored for classification.

language that the early layers of an encoder capture lower-level syntactical features better than the later layers [6, 24, 32, 40], which can benefit downstream tasks.

Inspired by the line of research that exploits early layers of models, we propose EarlyBIRD,² a novel and generic approach for building composite representations from the early layers of a pre-trained encoder model. EarlyBIRD aims to leverage all available information in existing pre-trained encoder models during fine-tuning to either improve results or achieve competitive results at reduced resource usage during code classification. We empirically evaluate EarlyBIRD on CodeBERT [14], a popular pre-trained encoder model for code, and four benchmark datasets that cover three common SE tasks: defect detection with the Devign and ReVeal datasets [20, 51], bug type inference with the data from Yasunaga et al. [47], and exception type classification [7]. The evaluation compares the *baseline* representation that uses the last encoder layer with results obtained via EarlyBIRD. We both fine-tune the full-size encoder and its pruned version with only several early layers present in the model. The latter scenario analyzes the trade-off between only using a partial model and the performance impact on SE tasks.

Contributions: In this paper, we make the following contributions: (1) *We propose EarlyBIRD, an approach for creating composite representations of code using the early layers of a transformer-based encoder model.* The goal is to achieve better code classification performance at equal resource usage or comparable performance at lower resource usage.

(2) *We conduct a thorough empirical evaluation of the proposed approach.* We show the effect of using composite EarlyBIRD representations while fine-tuning the original-size CodeBERT model on four real-world code classification datasets. We run EarlyBIRD with 10 different random initializations of non-fixed trainable parameters and mark the EarlyBIRD representations that yield statistically significant improvement over the baseline.

(3) *We investigate resource usage and performance of pruned models.* We analyze the trade-off between removing the later layers of a model and the impact this has on classification performance.

Main findings: With EarlyBIRD, we achieve performance improvements over the baseline code representation with the majority of representations obtained from single early layers on the defect detection task and selected combinations on bug type and exception type classification. Moreover, out of the reduced-size models with pruned later layers, we obtain a +2 average accuracy improvement on Devign with 3.3x speed-up of fine-tuning, as well as +0.4 accuracy improvement with 3.7x speed-up on average for ReVeal.

The remainder of the paper is organized as follows. We present related work in Section 2 and provide background details of the study in Section 3. The methodology is described in Section 4 which is followed by experimental setup in Section 5. We present and discuss results in Section 6 and conclude with Section 7.

2 RELATED WORK

Here, we give an overview of language models for SE tasks and recent encoder models, specifically, as well as different approaches to use early layers of encoder models.

2.1 Transformers in Software Engineering

The availability of open source code and increased hardware capabilities popularized training and usage of Deep Learning, including NLP and Large Language Models (LLMs), for SE tasks. To date, deep NLP models have already been applied in at least 18 SE tasks [28]. Pre-trained language models available for fine-tuning on SE tasks largely build on the transformer architecture, sequence-to-sequence models, and the attention mechanism [8, 9, 42]. One widely used benchmark to test different deep learning architectures on SE tasks is CodeXGLUE [27]. The benchmark provides data, source code for model evaluation, and a leader-board ranking model performance on different tasks [27].

SE tasks can be translated to input sequence classification and generation of code or text. Examples of generative tasks in SE are code completion, code repair, generation of documentation from code and vice versa, and translation between different programming languages. Such tasks are frequently approached with neural machine translation models. Full transformer models for translation from a programming language (PL) to a natural language (NL) or PL-PL tasks include PLBART [1], PYMT5 [10], TFix [4], CodeT5 [43], Break-It-Fix-It [47]. Alternatively, generative models can include the decoder-only part of the transformer as in GPT-type models. In this case, the decoder both represents the input sequence and transforms it into the output sequence. Decoder-based models for code include, for example, Codex and CodeGPT [8, 27].

In the tasks that require code or documentation representation and their subsequent classification, the encoder-only architectures are used more frequently than in translation tasks. Examples of code classification problems are code clone detection, detection of general bugs, such as the presence of swapped operands, wrong variable names, syntax errors, or security vulnerabilities. A number of encoder models for code applied a widely-used bi-directional encoder, BERT [12], to pre-train it on code, with some modifications of the input. In this way, the CodeBERT [14], GraphCodeBERT [16], CuBERT [20], and PolyglotCodeBERT [2] models were created. In detail, the 12-layer RoBERTa-based CodeBERT model was pre-trained on NL-PL tasks in multiple PLs and utilized only the textual features of code. Note that RoBERTa is a type of BERT model with optimized hyper-parameters and pre-training procedures [26]. Together with the decoder-only CodeGPT model, the encoder-only CodeBERT model was used as a baseline in CodeXGLUE. GraphCodeBERT utilizes both textual and structural properties of code to encode its representations. PolyglotCodeBERT is the approach that improves fine-tuning of the CodeBERT model on a multi-lingual dataset for a target task even if the target task tests only one PL. This paper focuses on the fine-tuning strategies which, in contrast to PolyglotCodeBERT, do not increase the resource usage for fine-tuning. CuBERT is a 24-layer pre-trained transformer-based encoder tested on a number of code classification tasks, including exception type classification. We test the performance of the proposed EarlyBIRD composite representations on defect detection, including the use of one of CodeXGLUE benchmarks, as well as on error and exception type classification tasks. However, the goal of this paper is to achieve improvement over the baseline model when it is fine-tuned with composite code representations. We do not aim to compare results with other models, but rather propose

² Early-layer Based Improvement or Reduction of resources useD

an approach that is applicable to transformer-based encoders for source code and show its performance gains compared to the same model usage without the proposed approach.

2.2 Use of Early Encoder Layers

A number of studies explored different approaches to use information from early layers of DL models for sequence representation, such as probing single layers, pruning and variable learning rates. One way to leverage information from early model layers is to give different priority to layers while fine-tuning the models [19, 39]. For example, the layer-wise learning rate decay (LLRD) strategy and re-initialization of late encoder layers yielded improvement over the standard fine-tuning of BERT on NLP tasks [49]. The LLRD strategy was initially developed to tune the later encoder layers with larger learning rate. In this way, the later layers can be better adapted to a downstream task under consideration, because the later layers are assumed to learn complex task-specific features of input sequences [19]. Moreover, Peters et al. [33] showed that the performance of fine-tuning improves if the encoder layers are updated during fine-tuning in comparison with training only the classifier on top of fixed (frozen) encoder layers.

Pruning later layers of transformer models is another way to consider only early layers for fine-tuning [13, 31, 36]. Sajjad et al. [36] investigated how the performance of transformer models on NLP is affected when reducing their size by pruning layers. They considered six pruning strategies, including dropping from different directions, alternated layer dropping, or dropping layers based on importance, for four pre-trained models: BERT [12], RoBERTa [26], XLNET [46], ALBERT [22]. By pruning model layers, Sajjad et al. were able to reduce the number of parameters to 60% of the initial parameter set while maintaining a high level of performance. While the performance on downstream tasks varies in their study, the lower layers are critical for maintaining performance when fine-tuning for downstream tasks. In other words, dropping earlier layers is detrimental to performance. Overall, pruning layers reduces model size and in turn reduces fine-tuning and inference time. In line with the work of Sajjad et al. [36], we extend our experiments with the pruning of later layers and keeping earlier layers present in the model (see RQ2 in Section 6).

The use of information from single early layers in a number of EarlyBIRD experiments is also inspired by Peters et al. [32]. In their study, Peters et al. present an empirical evidence that language models learn syntax and part-of-speech information on earlier layers of a neural network, while more complex information, such as semantics and co-reference relationships, are captured better by deeper (later) layers. In another study, Karmakar and Robbes probed pre-trained models of code, including CodeBERT, on tasks of understanding syntactic information, structure complexity, code length, and semantic information [21]. While Karmakar and Robbes probed frozen early layers of different models for code in a single strategy, we use 12 different strategies for combining unfrozen early layers during fine-tuning and focus on the tasks of bug detection or bug type classification. Similarly, Hernández López et al. [18] probed different layers of five pre-trained models, including CodeBERT [14] and GraphCodeBERT [16], and found that most syntactic information is encoded in the middle layers. The novelty of our

study with respect to Karmakar and Robbes is that we combine early layers in addition to extracting each of them, while Karmakar and Robbes extracted early layer representations and used them without composing new representations.

3 ENCODERS FOR CODE CLASSIFICATION

In this section, we present the background on transformer models and different uses of the encoder-decoder—or full transformer—architecture, as well as its encoder-only and decoder-only variants. Because our study focuses on encoder-only open-source models available for fine-tuning, the distinction between transformer types is necessary for understanding the methodology.

In sequence-to-sequence generation scenarios, the transformer model consists of a multi-layer encoder that represents the input sequence and a decoder that generates the output sequence based on the sequence representation from the encoder and the available output generated at previous steps [42]. For source code classification tasks, the transformer is frequently reduced to only its encoder followed by a *classification head*, a component added to the encoder to categorize the representation into different classes. Dropping the decoder for classification is motivated by resource efficiency, because the decoder is conceptually only needed for token generation from the input sequence. During classification of an input, the encoder represents the sequence and passes it to the classification head. Based on this design, a number of pre-trained encoders have been published in recent years, such as BERT and RoBERTa which were pre-trained on natural language, and similar models pre-trained on code, or a combination of code and natural language [12, 26]. The goal of pre-training in the *pre-train and fine-tune* scenario is to capture language patterns in general, so that they can serve as a basis for domain-specific downstream tasks. Pre-trained models can be fine-tuned on different downstream tasks in NLP and SE.

Processing the input sequence for classification consists of several steps: *tokenization*, *initial embedding*, *encoding* the sequence with an encoder, and passing the sequence representation through a *classification head*. Tokenization splits the input sequence, adds special tokens, matches the tokens to their ID's in the vocabulary of tokens, and unifies the resulting token length for samples in a dataset. Embedding transforms the one-dimensional token ID to an initial multi-dimensional static vector representation of the token and is usually a part of the pre-trained encoder model. This representation is updated using the attention mechanism of the encoder. Because of attention, the representation of the input is influenced by all tokens in the sequence, so it is contextualized.

CodeBERT is a RoBERTa-based model with 12 encoder layers pre-trained on 6 programming languages (Python, Java, JavaScript, PHP, Ruby, and Go), as well as text-to-code tasks [14]. Pre-training was done on the masked language modeling (MLM) and replaced token detection (RTD) tasks. These tasks respectively train the model to derive what token is masked in MLM, and in RTD predict whether any token in an original sequence is swapped with a different token that should not be in the sequence. CodeBERT outputs a bidirectional encoder representation of the input sequence, which means that the model considers context from pre-pending and subsequent words to represent each token in the input sequence.

A pre-trained model is usually released with a pre-trained tokenizer. The pre-trained tokenizer ensures that token ID's correspond to those processed during pre-training. The tokenizer also adds special tokens, such as a CLS token at the start of each input sequence, PAD tokens to unify lengths of input sequences, and the EOS token to signify the end of the input string and the start of padding sequence [12]. All tokens are transformed by the model in each encoder layer. Out of all tokens, the CLS token representation from the last layer, which is updated by all encoder layers, is typically used as a representation for the full sequence.

The standard practice of using the CLS token from the last encoder layer is motivated by the pre-training procedure. For example, in MLM, the model predicts the masked token based on the CLS token representation from the 12th layer of BERT and CodeBERT. However, the choice of token to represent the full sequence in fine-tuning can be different. For example, in PLBART [1], a transformer model for code with both an encoder and a decoder, the EOS token is used for representing the input sequence. In this paper, we propose different ways to represent the input sequence and use information from early layers of the model in an effective way.

4 METHODOLOGY

In this paper, the architecture of the code classification model consists of five parts: (1) a tokenizer, (2) an embedding layer, (3) an encoder with several layers, (4) a set of operations to combine sequence representations from encoder layers with EarlyBIRD, and (5) a classification head. The output of each step is used as input into the next step. An overview of the architecture is shown in Figure 1 and described below. The main difference between this architecture and the classification architecture discussed in Section 3 is step (4); the standard architecture only consists of steps (1–3) and (5).

Steps (1)–(3) use a pre-trained tokenizer, embedder, and encoder. EarlyBIRD is formulated in a generic way and can be applied to any encoder, but for our experiments, we fix the CodeBERT model and tokenizer. In step (4), we combine information from all the layers or from only some of the early layers of the encoder, as opposed to the baseline that uses the last layer of the encoder. Finally, the classification head in step (5) consists of one dropout layer and one linear layer with softmax.

The encoder model represents each token of an input sequence with a vector of size H , also known as hidden size. For each input sequence of length S , and a hidden size H , we obtain a matrix of size $S \times H$ for each of L layers of the base model as shown in Figure 1.

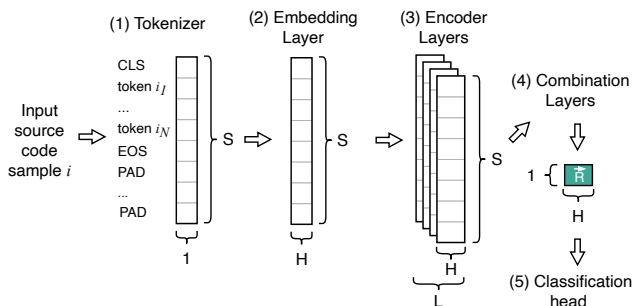


Figure 1: Model architecture for code classification.

For example, the CodeBERT architecture is fixed with 12 encoder layers, i.e., $L = 12$ for that model. All the information available in the encoder for one input sequence is stored in a tensor of size $L \times S \times H$. The EarlyBIRD combinations must produce one vector \vec{R} of size H that represents the input, as shown in Figure 1. Keeping the output code representation of size H is required to provide a fair comparison of EarlyBIRD composite representations with the standard code representation obtained from the last layer. In this way, the dimension of the classification head is the same for all combinations of early layers and has minimal possible influence during fine-tuning.

As a strategy for systematically investigating composite representations, we create a grid-search over three typical operations to combine outputs of neural network layers – maximum pooling (max pool), weighted sum and slicing – and two dimensions to apply the operations: over tokens and/or layers. For the tokens dimension, we either use all of the tokens from a specific layer or only the CLS token. Among layers, we either slice one layer, sum or take maximum values over all layers. The choice of considering every token of a layer is motivated by the fact the transformer-based models exhibit varying degrees of attention for different types of tokens [29], which indicates that solely using the CLS token might not be the best choice for tasks [37]. We also experiment with different sizes of the model. The combination strategies that use all layers of the pre-trained model are divided into two categories: the strategies that use CLS tokens from the encoder layers; the strategies that use more tokens than just CLS from encoder layers.

When we slice the CLS token and apply each of the operations over layers, we obtain the following CLS-token combinations:

- (i) *baseline*: CLS token from the last layer, i.e., layer no. L ;
- (ii) CLS token from one layer³ no. l , $l \in \{1, \dots, (L - 1)\}$;
- (iii) max pool over CLS tokens from all layers $\{l\}_{l=1}^L$;
- (iv) weighted sum over CLS tokens from all layers $\{l\}_{l=1}^L$.

The second set of combinations uses representations of all the tokens in tokenized input sequences, including the CLS token. We first apply max pooling operation to either all tokens or all layers and use the rest of operations. Then we apply weighted sum as the first operation followed by max pool or slicing of a layer:

- (v) max pool tokens from one layer no. l , $l \in \{1, \dots, L\}$;
- (vi) max pool over all layers for each token in the input sequence, max pool over tokens;
- (vii) max pool over all layers for each token in the input sequence; weighted sum over tokens;
- (viii) max pool over all tokens for each layer no. l , $l \in \{1, \dots, L\}$; weighted sum over layers
- (ix) weighted sum over tokens from one layer no. l , $l \in \{1, \dots, L\}$;
- (x) weighted sum over tokens for each one layer no. l , $l \in \{1, \dots, L\}$; weighted sum over all layers;
- (xi) weighted sum over all layers for each token in the input sequence; weighted sum over all tokens.

Note that weights in the weighted sums are learnable parameters. However, the added number of learnable parameters for fine-tuning

³ We use each layer l in the combinations separately if we denote l , $l \in \{1, \dots, L\}$, and specify the set of layers $\{l\}_{l=1}^L$ if several layers are used at once.

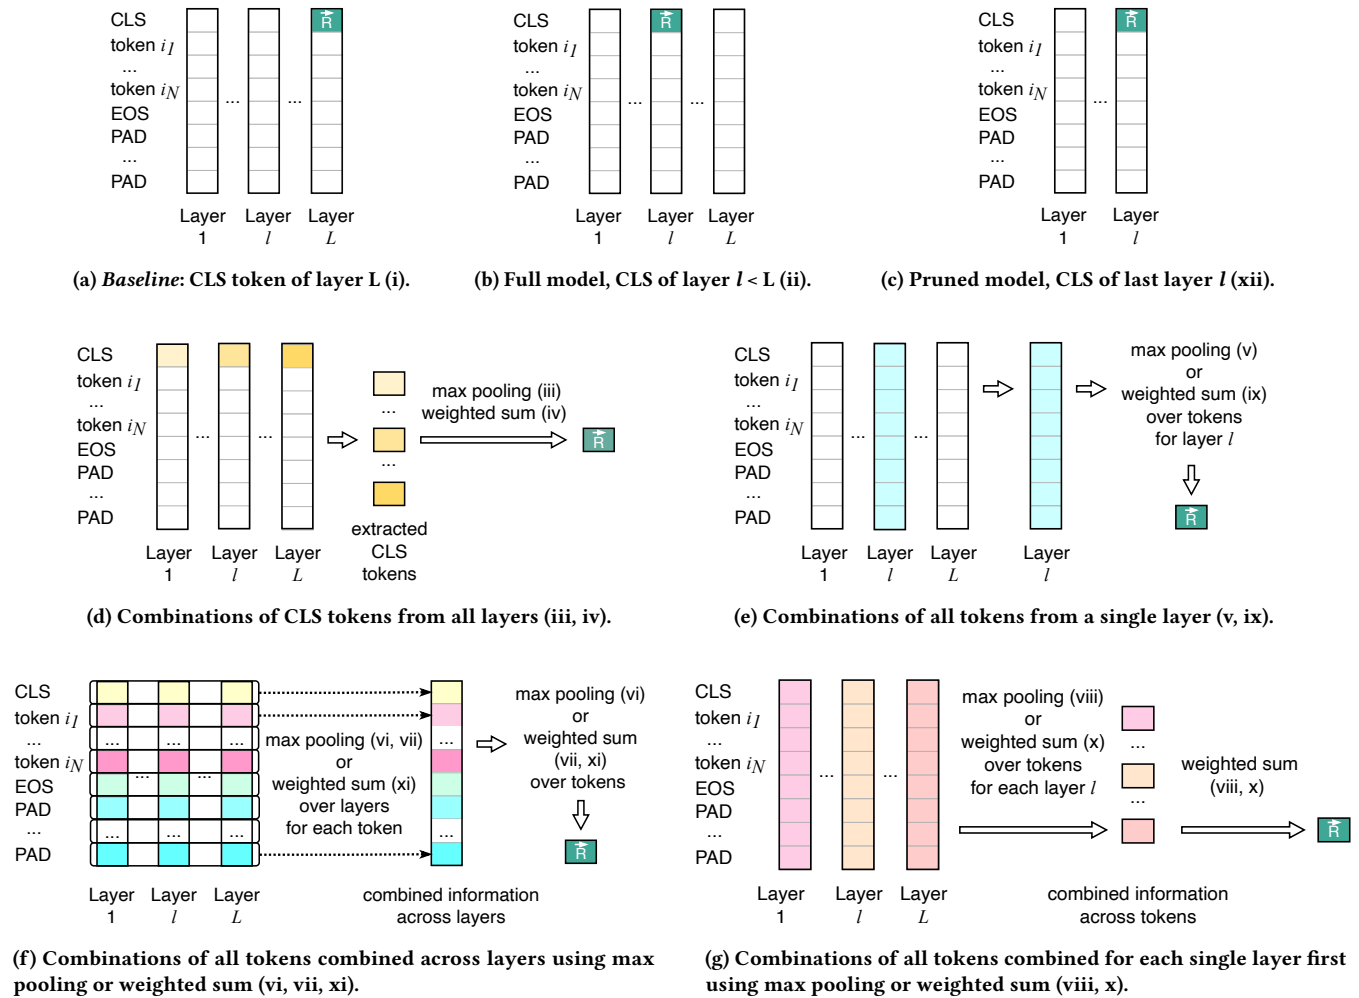


Figure 2: Combinations of early encoder layers that lead to code representation vector \vec{R} for each tokenized input sequence. The latin numbering in brackets corresponds to the combinations described in Section 4. Observe that the presentation order has been designed to preserve space by grouping similar combinations in the same subfigure.

constitutes 0.00042%⁴ of the number of learnable parameters in the baseline configuration. For this reason, we mention that the models with combinations (ii-x) have the same model size while bearing in mind the overhead of learnable weights in the weighted sums.

In addition to experiments with token combinations, we also investigate performance of the model with first $l < L$ layers and the baseline token combination, described as follows:

- (xii) CLS token from the last layer of the model with $l < L$ encoder layers.

Note that the baseline combination (i) with the usage of the CLS token from layer L corresponds to (ii) and (xii) if $l = L$.

⁴ Weighted sum over tokens adds $S = 512$ learnable weights. Because the weights of the sum are shared across the layers, the maximum number of added weights is $L + H = 524$ out of 124M learnable weights in the base model. Combinations without weighted sums do not add extra learnable parameters to the base model. Weighted sum over layers adds learnable $L = 12$ weights for CodeBERT.

The combinations are presented in Figure 2. Similar combinations are presented close to each other or are combined in the same image if they only have minor differences and share the major parts. For example, in Figure 2c, we illustrate combinations (iii) and (iv), because both of them use CLS tokens from all layers combined using max pooling or weighted sum. The roman numbers which indicate combination types are preserved either in the descriptions below the figures or in the figures themselves, but the order is changed. We mention combination number corresponding to the description in the current section, such as baseline combination (i) in Figure 2a or combination (ii) for CLS token from one early layer in Figure 2b. We highlight what parts of encoder layer outputs are used for each combination with color. White cells correspond to the tokens that are not used in early layer combinations. The goal of all combinations is to obtain a vector representation \vec{R} for each input code

sample. For example, in Figure 2a, we consider the last layer L and extract only the CLS token marked as \vec{R} .

Another remark on the EarlyBIRD combinations concerns the usage of all tokens or only code tokens. Code tokens are those that correspond to tokenized input words or sub-words and are shown in Figure 2 as $\text{token}_{i_1}, \dots, \text{token}_{i_N}$ for an input sequence i of size i_N . For each combination that uses more than just a CLS token, i.e., combinations (v-xi), we experiment with code tokens only, as well as with all tokens, including CLS, EOS, and PAD. The motivation to check code tokens exclusively stems from the hypothesis that information in special tokens may introduce noise into results.

5 EXPERIMENTAL SETUP

In this section, we describe the datasets used for empirical evaluation and implementation details of fine-tuning with the proposed EarlyBIRD approach. We investigate binary and multi-task code classification scenarios to explore generalisability of our results.

5.1 Datasets for Source Code Classification

We fine-tune and test the CodeBERT model using the EarlyBIRD approach on four datasets. The datasets span three tasks: defect detection, error type classification and exception type classification – with 2, 3, and 20 classes, respectively. They also contain data in two programming languages, C++ and Python. In addition, the chosen datasets have similar train subset sizes. In this way, we aim to reduce the effect of the model’s exposure to different amounts of training data during fine-tuning. Statistics of the datasets are provided in Table 1. We report the size of the train/validation/test splits. In addition, we compute the average number of tokens in the input sequences upon tokenization with the pre-trained CodeBERT tokenizer. Because the maximum input sequence size for the CodeBERT model is limited to $S = 512$, the number of tokens is indicative of how much information the model gets access to or how much information is cut off, in case of long inputs.

Devign: This dataset contains functions in C/C++ from two open-source projects labelled as vulnerable or non-vulnerable [51]. We reuse the train/validation/test split from the CodeXGLUE Defect detection benchmark.⁵ The dataset is balanced: the ratio of non-vulnerable functions is 54%.

ReVeal: Similarly to Devign, ReVeal is a vulnerability detection dataset of C/C++ functions [7]. The dataset is not balanced: it contains 90% non-vulnerable code snippets. Both the Devign and ReVeal datasets contain real-world vulnerable and non-vulnerable functions from open-source projects.

Break-It-Fix-It (BIFI): The dataset contains function-level code snippets in Python with syntax errors [47]. We use the original buggy functions and formulate a task of classifying the code into three classes: Unbalanced Parentheses with 43% of the total number of code examples in BIFI, Indentation Error with 31% code samples, Invalid Syntax containing 26% samples. The train/test split provided in the dataset is reused, and the validation set is extracted as 10% of training data.

Exception Type: The dataset consists of short functions in Python with an inserted `__HOLE__` token in place of one exception in code.⁶

⁵ <https://github.com/microsoft/CodeXGLUE/tree/main/Code-Code/Defect-detection>
⁶ <https://github.com/google-research/google-research/tree/master/cubert>

Table 1: Statistics of Fine-Tuning Datasets.

Dataset	# classes	Avg # tokens	# code samples		
			Train	Valid	Test
Devign	2	614	21,854	2,732	2,732
ReVeal	2	512	18,187	2,273	2,274
BIFI	3	119	20,325	2,259	15,055
Exception Type	20	404	18,480	2,088	10,348

The task is to predict one of 20 masked exception types for each input function and is unbalanced. The dataset was initially created from the ETH Py150 Open corpus⁷ as described in the original paper [20]. We reuse the train/validation/test split provided by the authors.

5.2 Implementation

The architecture is based on the CodeBERT⁸ tokenizer and encoder model. The model defines the maximum sequence length, hidden size, and has 12 layers, so $S = 512$, $H = 768$, $L = 12$. Hyperparameters in the experiments are set to $B = 64$, learning rate is $1e-5$, and dropout probability is 0.1. If the tokenized input sample is longer than $S = 512$, we prune the tokens in the end to make the input fit into the model. We run fine-tuning with Adam optimizer and testing for each combination 10 times with different seeds for 10 epochs and report the performance for the best epoch on average over 10 runs. The best epoch is defined by measuring accuracy on a validation set. We use Python 3.7 and Cuda 11.6, and run experiments on one Nvidia Volta A100 GPU.

5.3 Evaluation Metrics

To present the impact of early layer combinations, we compare the accuracy on the test set for all datasets, because it allows us to compare our results with other benchmarks. In addition, we report weighted F1-score denoted as $F1(w)$ for a detailed analysis of selected combinations to account for class imbalance. To obtain the weighted F1-score, the regular F1-score is calculated for each label and their weighted mean is taken. The weights are equal to the number of samples in a class.

We also report results of the Wilcoxon signed-rank test on the corresponding metrics for the combinations that show improvement over the baseline [45]. The Wilcoxon test is a non-parametric test suitable for the setting in which different model variants are tested on the same test set, because it is a paired test. The Wilcoxon test checks the null hypothesis whether two related paired samples come from the same distribution. We reject the null hypothesis if p-value is less than $\alpha = 0.05$. In case we obtain improvement of a metric over the baseline with an EarlyBIRD combination and the null hypothesis is rejected, we conclude that the combination performs better and the result is statistically significant. For the pruned models, we compute Vargha and Delaney’s A_{12} non-parametric effect size measure of the performance change for accuracy and $F1(w)$

⁷ <https://www.sri.inf.ethz.ch/py150>

⁸ <https://huggingface.co/microsoft/codebert-base>

with thresholds of 0.71, 0.64 and 0.56 for large, medium and small effect sizes [41].

5.4 Research Questions

During our empirical evaluation of composite EarlyBIRD code representations, we address the following research questions:

RQ1. Composite Code Representations with Same Model Size: What is the effect of using combinations (ii-xi) of early layers with the same model size in comparison to the baseline approach of using only the CLS token from the last layer, i.e., combination (i), for code representation on model performance in the code classification scenario? The goal is to find out whether any of the EarlyBIRD combination types work consistently better for different datasets and tasks.

RQ2. Pruned Models: What is the effect of reducing the number of pre-trained encoder layers in combinations (xii) on resource usage and model performance on code classification tasks? As opposed to RQ1, in which we consider the combinations that do not reduce the model size, this research question is devoted to investigation of the trade-off between using less resources with reduced-size models and performance variation in terms of classification metrics.

For both research questions, we evaluate the composite representations on binary and multi-task code classification scenarios to explore generalisability of the results obtained for the binary case. We investigate if and what combinations result in better performance, averaged over 10 runs with different seeds. For combinations that improve the baseline on average, we also explore if the results are statistically significant according to the Wilcoxon test.

6 RESULTS AND DISCUSSION

6.1 EarlyBIRD with Fixed-Size Models

To answer RQ1, we explore one-layer combinations, multi-layer combinations, and estimate the statistical significance of the performance improvement.

6.1.1 Combinations of Tokens in Single Selected Early Layers. Figure 3 shows a heatmap of the difference of the mean accuracy obtained with each combination that uses only one selected early layer compared to the baseline. In addition, we show the value of the difference in mean accuracy for each combination type and layer number. Note that the scale is logarithmic and in the most extreme case spans the interval from ca. -37 to +2. Negative values are shown in black, and positive values are shown in white. Differences that are statistically significant according to the Wilcoxon test are marked with a star (*) next to the value. Combinations that correspond to the baseline are marked with “bsln” and have zero difference, by definition. The results for the weighted F1-score show a similar pattern as those for the mean accuracy. They are visualized in the same way in Figure 4.

The first rows in Figures 3a and 3b correspond to the combinations (ii) CLS token layer *l*. With this combination type, average improvement over the baseline is achieved with the majority of early layers. Specifically, we obtain accuracy improvements ranging from +0.2 to +2.0 for Devign in 8 out of 11 layers, and accuracy improvements from +0.1 to +0.8 for ReVeal in 9 out of 11 layers.

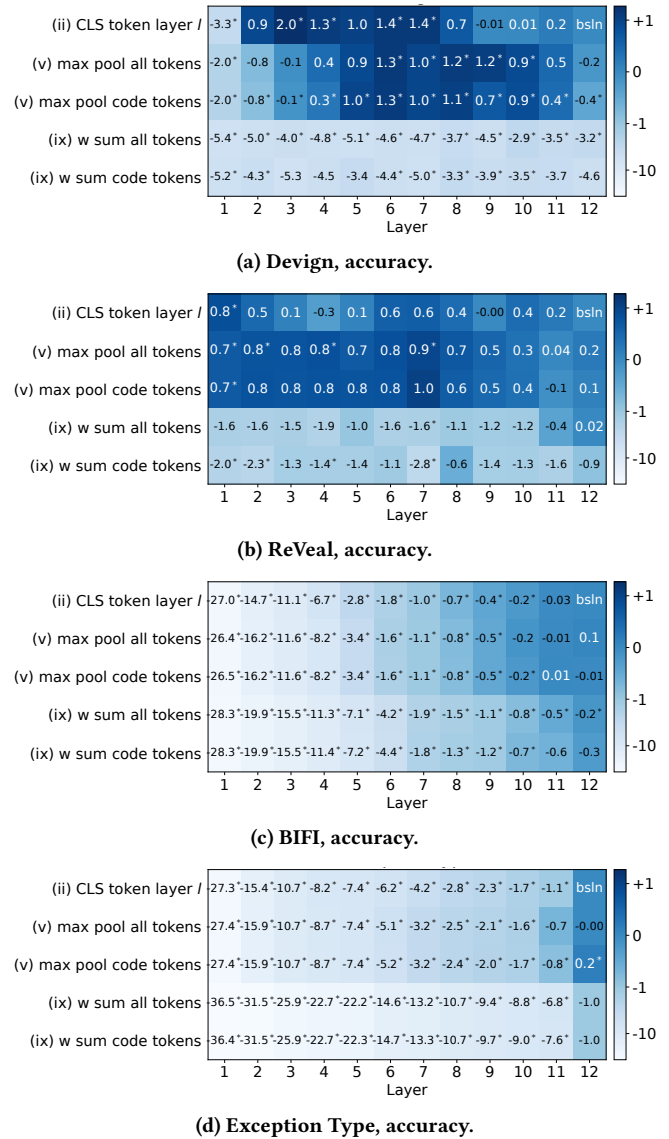


Figure 3: Difference of mean accuracy between EarlyBIRD and baseline (bsln) performance. The star * indicates a statistically significant difference w.r.t. the baseline.

The dynamic of the metric change over selected layer numbers is different for Devign and ReVeal. In detail, the average performance of combination (ii) is best with layer 3 on Devign (a +2.0 accuracy improvement) and with layer 1 for ReVeal (a +0.8 accuracy improvement). The best improvement in terms of F1(w) matches with layer 3 for Devign and with layer 2 for ReVeal, as shown in Figure 4.

Max pooling over all available tokens from a selected layer in combination (v) also achieves performance improvement over the baseline, as shown in rows 2 and 3 of Figures 3a, 3b. In general, layers 4–11 yield higher accuracy and layers 2–11 higher F1(w) with max pooling for Devign than the baseline. For ReVeal, all layers except layer 11 result in better average accuracy and layers 2–10

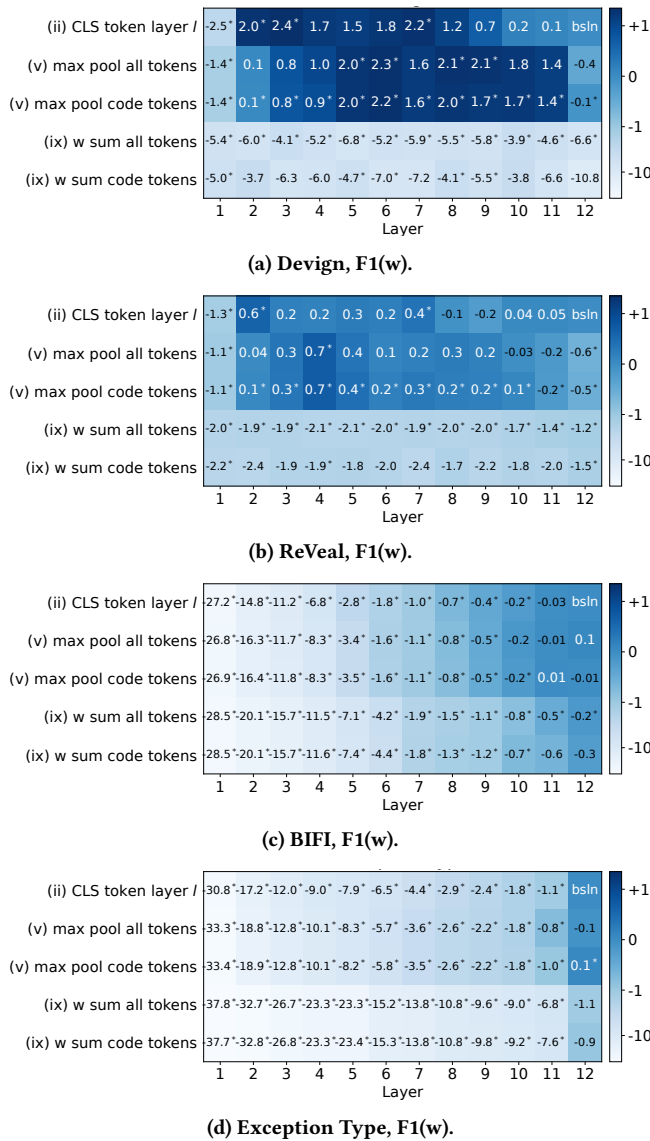


Figure 4: Difference of mean weighted F1-scores (F1(w)) between EarlyBIRD and baseline (bsln). The star * indicates a statistically significant difference w.r.t. the baseline.

have higher average F1(w). Max pooling over all tokens, including special tokens, achieves the best statistically significant average improvement of accuracy of +0.9 of all combinations for ReVeal.

The weighted sum of all tokens or code tokens exclusively in combination (ix) does not improve the baseline performance. We assume that fine-tuning for 10 epochs is not enough for this type of combination, because the loss at epoch 10 on both training and validation splits is higher for combinations (ix) than for combinations with max pooling. Since the goal of this study is to use the same or less resources for fine-tuning, we have not fine-tuned this combination for more than 10 epochs.

While combinations (ii) and (v) perform better for the majority of layers on the defect detection task, multi-class classification for bug or exception type prediction does not benefit from the combinations to the same extent as the binary task. Only max pooling of tokens of the last encoder layer achieves better performance than the baseline for BIFI (+0.1 accuracy, +0.1 weighted F1-score improvements) and Exception Type (a +0.2 accuracy, +0.1 weighted F1-score improvements) datasets.

The impact of using all tokens or code tokens exclusively depends on the dataset. The difference between performance of single-layer combinations with max pooling of all tokens and only code tokens constitute 0.0-0.1 accuracy or F1(w). For the multi-class tasks, the average results improve with the use of each later layer in the model. We obtain performance improvement with the max pooling combination (v), while other one-layer combinations do not perform better than the baseline.

The best performing results on Devign and Exception Type classification datasets are statistically significant according to the Wilcoxon test. For ReVeal, the second best result is statistically significant improvements for BIFI. We explain it by the fact that the baseline metric is already high, i.e., 96.7 accuracy. Achieving improvement is usually more challenging when the baseline performs at this level.

In essence, the combinations that involve CLS tokens corresponding to the single layer (ii), as well as the max pooling combinations (v) perform better on average for defect detection datasets Devign and Reveal. However, only the max pooling combination (v) of tokens from the last encoder layer outperforms the baseline on average for multi-class datasets BIFI and Exception Type. The weighted sum of tokens from a selected layer (ix) performs worse than the baseline if fine-tuned for the same number of epochs for all tasks. Multi-class classification tasks require the information from the last layer for better performance in our experiments, while the binary task of defect detection allows us to use early layers and improve the performance over the baseline.

6.1.2 Multi-Layer Combinations. The average performance difference with the baseline of combinations that utilize early layers is shown as heatmaps in Figures 5 and 6. We include the value of the average performance difference and add a star (*) to the number if the difference is statistically significant. Again, negative values are shown in black, and positive values are shown in white.

When we use all information from the available layers, the improvement over the baseline is less than what was observed in Section 6.1.1, where one specific layer has been used. In detail, out of combinations that involve CLS tokens from all early layers, no combination performs better than the baseline for ReVeal, BIFI, or Exception Type datasets. However, the best improvement (+0.6 accuracy) out of experiments with all layers is obtained on Devign with the weighted sum of CLS tokens in the combination (iv), which is less than the maximum improvement with the combinations from one selected early layer in Section 6.1.1. The improvement of F1(w) is shown in Figure 6. We obtained slightly better improvements of F1(w) for Devign, no F1(w) improvement for the unbalanced ReVeal dataset. The average F1(w) difference with the baseline for multi-class tasks are the same as accuracy difference.

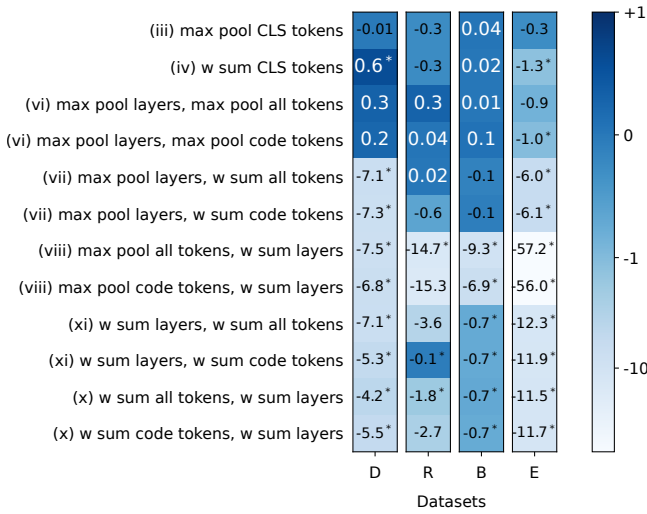


Figure 5: Difference of average accuracy between EarlyBIRD and baseline performance on D (Devign), R (ReVeal), B (BIFI), E (Exception Type). The star * indicates a statistically significant difference w.r.t. the baseline.

If we consider the combinations that involve all tokens, the combination (vi) with two max pooling operations outperforms the baseline for Devign, Reveal, and BIFI with accuracy improvement between +0.1 and +0.3. No combination that involves all layers outperforms the baseline on average for Exception Type dataset. Combinations that involve one max pooling and one weighted sum of all tokens perform worse or neutral in comparison with the baseline. The combinations with only weighted sums perform worse than the baseline on average.

Answer to RQ1. EarlyBIRD achieves statistically significant accuracy and F1-score improvements for defect detection datasets by using single-layer combinations that involve the CLS token or max pooling over all tokens. For bug type and exception type classification, max pooling of the tokens from the last encoder layer has improved the performance. Weighted sum of tokens does not improve performance over the baseline.

6.2 Pruned Models

This section is devoted to the combinations of early layers that are initialized with the first $l < L$ early layers from the pre-trained model and fine-tuned as l -layer models – combinations (xii). We start by comparing the performance of using the CLS token from layer l of the full-size model, i.e., combination (ii), and using the CLS token from layer l of the model that has l layers in total – combination (xii). Figure 7 presents average accuracy obtained with these two combinations depending on the used layer, as well as the baseline combination of using CLS from the last layer $L = 12$ of CodeBERT. On average, the pruned models with reduced size perform on par with the full-size model for defect detection on the balanced Devign dataset, and for bug type and exception type classification. However, the performance of the two analogical

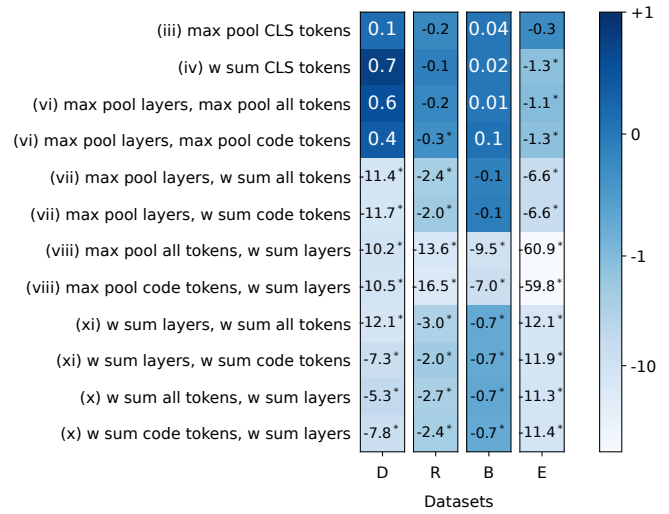


Figure 6: Difference of mean weighted F1-score between combinations and baseline performance. Datasets are abbreviated to D (Devign), R (ReVeal), B (BIFI), E (Exception Type). The star * indicates a statistically significant difference w.r.t. the baseline.

combinations diverges for the unbalanced defect detection dataset ReVeal in layers 4 and 6–11.

Most importantly, the results show that reducing the model size and using the CLS token from the last layer of the reduced model performs on par with the baseline for the defect detection task. The best improvement with the reduced model is achieved with the 3-layer encoder for Devign and the 1-layer encoder for ReVeal. This result shows that it is possible to both reduce resources and improve the model’s performance during fine-tuning on the defect detection task with both a balanced and unbalanced dataset.

To explore the trade-off between resource usage and performance degradation for bug type and exception type identification,

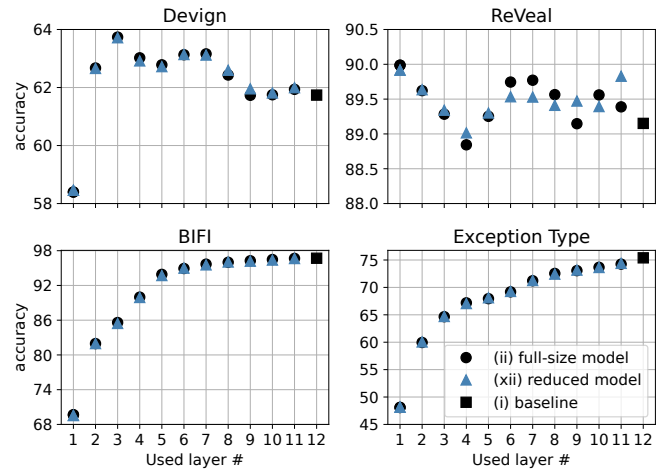


Figure 7: Model performance with a subset of $l < L$ layers (xii) vs. models with all layers (ii); CLS token from layer l .

Table 2: Comparison of reduced size models with the baseline. We report metric performance for the baseline and difference with the baseline for reduced models, average time for one-epoch fine-tuning (Time) in mm:ss format, speed-up and performance variation obtained with models with l layers. Statistically significant improvements are marked in bold, statistically insignificant performance losses are marked with *, and A_{12} effect sizes for accuracy and F1(w), if any, are indicated by cell color, respectively large, medium, and small. The best metric improvement with highest speed-up factors are underlined.

l	Devign				ReVeal				BIFI				Exception Type			
	Time	Speed-up	Acc	F1(w)	Time	Speed-up	Acc	F1(w)	Time	Speed-up	Acc	F1(w)	Time	Speed-up	Acc	F1(w)
12	8:50	1.0x	61.7	60.4	6:57	1.0x	89.2	88.5	8:41	1.0x	96.7	96.7	7:22	1.0x	75.4	75.3
11.0	8:03	1.1x	+0.3	-0.1*	6:56	1.0x	+0.6	+0.3	7:07	1.2x	-0.1*	-0.1*	6:33	1.1x	-1.0	-1.0
10.0	7:13	1.2x	+0.1	+0.3	6:20	1.1x	+0.2	-0.0*	6:22	1.4x	-0.3	-0.3	6:01	1.2x	-1.8	-1.8
9.0	6:40	1.3x	+0.3	+0.5	5:53	1.2x	+0.3	-0.1*	5:46	1.5x	-0.5	-0.5	5:29	1.3x	-2.2	-2.3
8.0	5:52	1.5x	+0.9	+1.1	5:15	1.3x	+0.2	-0.1*	5:13	1.7x	-0.6	-0.6	4:55	1.5x	-3.0	-3.0
7.0	5:23	1.6x	+1.4	+2.2	4:44	1.5x	+0.3	+0.2	4:43	1.8x	-1.1	-1.1	4:20	1.7x	-4.1	-4.4
6.0	4:54	1.8x	+1.4	+2.0	4:25	1.6x	+0.3	+0.1	4:10	2.1x	-1.7	-1.7	3:50	1.9x	-6.1	-6.6
5.0	4:03	2.2x	+1.0	+1.5	3:45	1.9x	+0.1	+0.2	3:31	2.5x	-3.0	-3.0	3:15	2.3x	-7.3	-7.7
4.0	3:22	2.6x	+1.2	+1.6	3:06	2.2x	-0.2*	+0.2	2:53	3.0x	-6.8	-6.8	2:41	2.7x	-8.3	-9.2
3.0	2:41	<u>3.3x</u>	+2.0	+2.4	2:28	2.8x	+0.1	+0.3	2:15	3.8x	-11.3	-11.3	2:10	3.4x	-10.7	-12.0
2.0	2:00	4.4x	+1.0	+1.9	1:52	<u>3.7x</u>	+0.4	+0.6	1:34	5.5x	-14.7	-14.8	1:34	4.7x	-15.4	-17.1
1.0	1:18	6.8x	-3.2	-2.3	1:13	5.7x	+0.8	-1.2	0:58	9.0x	-27.2	-27.3	0:59	7.4x	-27.3	-30.7

we show the average speed-up of one fine-tuning epoch and the performance loss compared to the baseline for BIFI and Exception Type datasets in Table 2. We also report the corresponding values for Devign and ReVeal, for which both gains and losses of performance are indicated. The speed up is reported as a scaling factor of the baseline time. The metric difference is shown as gain or loss of the weighted F1-score and accuracy compared to the baseline performance. Statistically significant improvements are reported in bold, while statistically insignificant losses are marked with a star (*). The A_{12} effect sizes are indicated by three shades of blue as the cell color, with the darkest shade indicating a large effect ($A_{12} > 0.71$), the middle shade indicating a medium effect ($A_{12} > 0.64$), and the lightest shade indicating a small effect ($A_{12} > 0.56$). We also underline and discuss selected results that improve the metric values and reduce resource usage.

The majority of combinations (xii) with pruned models outperform the baseline for Devign and ReVeal. Furthermore, models with 2–10 layers show statistically significant improvements of both metrics on Devign, with the 3-layer model achieving +2 accuracy improvement with a 3.3-times average speed-up of fine-tuning with the same hardware and software. Not only does the 3-layer model improve the accuracy over CodeBERT baseline to 63.7, but also outperforms several other models tested on Devign and reported on the CodeXGLUE benchmark [27]. In particular, our pruned 3-layer CodeBERT model outperforms the full-transformer model PLBART [1], and code2vec code representations pre-trained on abstract syntax trees and code tokens in a joint manner [1]. However, our pruned model does not outperform the best performing model reported on CodeXGLUE, CoText, which achieves 66.62 accuracy [34].

Models with 1 and 11 layers achieve statistically significant accuracy improvements for ReVeal. However, the 1-layer model reduces the F1(w) score. The use of layer 11 does not impact the speed of fine-tuning, while the 1-layer model yields the 3.7x acceleration of the baseline fine-tuning speed. The lack of speed-up with 11-layer

model can be explained by the fact that the number of trainable parameters does not decrease linearly with the removal of later layers, since the additional embedding layer and classification head remain unchanged. The 2-layer model results in the best improvement of F1(w) which is statistically significant. The 2-layer model improves accuracy on ReVeal as well. For Devign and ReVeal, statistically significant improvements have large effect size.

For BIFI, we obtain statistically insignificant decrease of F1(w) and accuracy according to the Wilcoxon test which brings about 1.2x speed-up of the fine-tuning with the 11-layer model. If we decrease the number of layers to 8, the performance on BIFI stays within the (baseline metric–1) limit, but we gain up to 1.7x average speed-up of one-epoch fine-tuning. In case of using models with 1–10 layers, we observe a statistically significant change of distribution and decrease of metric values.

For the unbalanced Exception Type dataset, the performance drops faster and the speed-up is less prominent than for BIFI. The change of mean values of the metrics for all models is statistically significant. In detail, the metrics decrease by -1.0 absolute metric value at 11 layers with 1.1x fine-tuning speed-up and by -1.8 with 10 layers with 1.2x speed-up. We explain the sharper decline of the combinations performance by the lower baseline metric values (75.39 accuracy, 75.30 weighted F1-score) than in the case of BIFI (96.7 accuracy and weighted F1-score). For BIFI, statistically insignificant deterioration have small effect size. However, for both BIFI and Exception Type datasets, we observe deterioration of performance of large effect size with pruned models.

We conclude that for the BIFI dataset with high-performing baseline and 3 classes, the performance loss at removing each layer is less than for the Exception Type classification dataset with lower baseline performance and 20 classes. The resource usage, which is correlated with time spent on tuning, decreases faster for BIFI than for Exception Type. This is partially explained by a larger classification head for the Exception Type dataset, because this dataset has 20 classes as opposed to only 3 classes in BIFI. In other

words, we observe that the BIFI dataset has a strong baseline that is hard to outperform with pruning. By contrast, the complexity of the Exception Type dataset can influence the results in the opposite way: The baseline performance is already not very strong, and it proves hard to further improve on it with early layers only.

Answer to RQ2. We obtain performance improvements over the baseline as well as fine-tuning speed-ups for both defect detection datasets by using the CLS token from the last layer of pruned models. For multi-class classification, performance decreases upon pruning each layer from the end of the model. The decrease is sharper for the dataset with 20 exception types than for the task with 3 bug types.

6.3 Threats to Validity

The main threat to external validity is that the results are empirical and may not generalize to all code classification settings, including other programming languages, tasks, and encoder-based models for code. We have tested EarlyBIRD combinations on code in C for defect detection and Python for bug type and exception type classification in this study. The choice of the CodeBERT as the encoder model and its internal structure affects the results. For instance, an encoder model that takes smaller input sequences can perform worse on the same datasets, because larger parts of input code sequences have to be pruned in this case. The external validity can be improved by testing on more datasets and encoder models.

The threats to internal validity concern the dependency of models on initializations of trainable parameters and the choice of methods. Classification head and weighted sums with trainable parameters in our experiments depend on the initialization of the parameters and can lead the model to arrive at different local minima during fine-tuning. To reduce the effect of different random initializations, we have fine-tuned and tested all EarlyBIRD combinations 10 times with different random seeds.

In addition, we used the Wilcoxon test to verify whether the achieved improvements are statistically significant. However, the Wilcoxon test only estimates whether measurements of baseline values and EarlyBIRD combinations are drawn from different distributions. The reported times spent on fine-tuning and corresponding speed-ups have the purpose of illustrating the reduction in resource usage, and will depend on the hardware used. Even when using factors of speed-up for pruned models, there is a chance that these numbers will be different on other hardware configurations.

We implemented the algorithms and statistical procedures in Python, with the help of widely used libraries such as PyTorch, NumPy and SciPy. However, we cannot guarantee the absence of implementation errors which may have affected our evaluation.

7 CONCLUDING REMARKS

In this paper, we have proposed EarlyBIRD, an approach to combine early layers of encoder models for code, and tested different early-layer combinations on the software engineering tasks of defect detection, bug type and exception type classification. Our study is motivated by the hypothesis that early layers contain valuable information that is discarded by the standard practice of representing the code with the CLS token from the last encoder layer.

EarlyBIRD provides ways to improve the performance of existing models with the same resource utilization, as well as for resource usage reduction while obtaining comparable results to the baseline. **Results:** Using EarlyBIRD, we obtain statistically significant improvements over the baseline for the majority of the combinations that involve a single encoder layer on defect detection, and with selected EarlyBIRD combinations on bug type and exception type classification. Max pooling of tokens from selected single layers yields performance improvements for all datasets. Both the classification performance and the average fine-tuning time for one epoch are improved by pruning the pre-trained model to its early layers and using the CLS token from the last layer of the pruned model. For defect detection, this results in a +2.0 increase in accuracy and a 3.3x fine-tuning speed-up on Devign, and up to +0.8 accuracy improvement with a 3.7x speed-up on ReVeal. Pruned models do not lead to multi-class classification performance gains, but they do show a fine-tuning speed-up and the associated decrease in resource consumption.

The results show that pruned models with reduced size either work better or can result in a reduction of resource usage during fine-tuning with different levels of performance variation, which indicates the potential of EarlyBIRD in resource-restricted scenarios of deploying defect detection and but type classification in production environments. For example, EarlyBIRD achieves a 2.1x speed-up for BIFI while reducing accuracy from 96.7 to 95.0.

Future Work: The study can be extended by investigating the generalization to other encoder models. We are in the process of studying the performance of EarlyBIRD with two new encoder models: StarEncoder [23] and ContraBERT_C [25]. Another direction for future research is whether the types of layer combination and pruning, as we have investigated in this paper for encoder architectures, are also effective techniques for decoder and encoder-decoder architectures. Moreover, it would be of interest to experiment with other code classification tasks, such as general bug detection and the prediction of vulnerability types. The latter could be investigated using the CWE types from the Common Weakness Enumeration as labeled in the CVEfixes dataset [5].

ACKNOWLEDGMENTS

This work is supported by the Research Council of Norway through the secureIT project (IKTPLUSS #288787). Max Hort is supported through the ERCIM ‘Alain Bensoussan’ Fellowship Programme. The empirical evaluation presented in this paper was performed on the Experimental Infrastructure for Exploration of Exascale Computing (eX3), financially supported by the Research Council of Norway under contract #270053, as well as on resources provided by Sigma2, the National Infrastructure for High Performance Computing and Data Storage in Norway.

DATA AVAILABILITY

To support open science and allow for reproduction and verification of our work, all artifacts are made available through Zenodo at the following URL: <https://doi.org/10.5281/zenodo.7608802>.

REFERENCES

- [1] Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified Pre-training for Program Understanding and Generation. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 2655–2668. <https://doi.org/10.18653/v1/2021.naacl-main.211>
- [2] Toufique Ahmed and Premkumar Devanbu. 2022. Multilingual Training for Software Engineering. In *Proceedings of the 44th International Conference on Software Engineering*. ACM, Pittsburgh Pennsylvania, 1443–1455. <https://doi.org/10.1145/3510003.3510049>
- [3] Miltiadis Allamanis, Earl T. Barr, Premkumar Devanbu, and Charles Sutton. 2018. A Survey of Machine Learning for Big Code and Naturalness. *Comput. Surveys* 51, 4 (July 2018), 81:1–81:37. <https://doi.org/10.1145/3212695>
- [4] Berkay Berabi, Jingxuan He, Veselin Raychev, and Martin Vechev. 2021. TFix: Learning to Fix Coding Errors with a Text-to-Text Transformer. In *International Conference on Machine Learning*, Vol. 139. PMLR, Virtual Event, 780–791.
- [5] Guru Bhandari, Amara Naseer, and Leon Moonen. 2021. CVEfixes: Automated Collection of Vulnerabilities and Their Fixes from Open-Source Software. In *International Conference on Predictive Models and Data Analytics in Software Engineering (PROMISE)*. ACM, 30–39. <https://doi.org/10.1145/3475960.3475985>
- [6] Tera Blevins, Omer Levy, and Luke Zettlemoyer. 2018. Deep RNNs Encode Soft Hierarchical Syntax. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Melbourne, Australia, 14–19. <https://doi.org/10.18653/v1/p18-2003>
- [7] Saikat Chakraborty, Rahul Krishna, Yangruibo Ding, and Baishakhi Ray. 2022. Deep Learning Based Vulnerability Detection: Are We There Yet? *IEEE Transactions on Software Engineering* 48, 9 (Sept. 2022), 3280–3296. <https://doi.org/10.1109/tse.2021.3087402>
- [8] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgun Guss, Alex Nichol, Alex Paino, Nikolai Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. <https://doi.org/10.48550/arXiv.2107.03374> arXiv:2107.03374 [cs]
- [9] Zimin Chen, Steve James Kommrusch, Michele Tufano, Louis-Noël Pouchet, Denys Poshyvanyk, and Martin Monperrus. 2019. SEQUENCER: Sequence-to-Sequence Learning for End-to-End Program Repair. *IEEE Transactions on Software Engineering* 47, 9 (2019), 1943–1959. <https://doi.org/10.1109/tse.2019.2940179>
- [10] Colin Clement, Dawn Drain, Jonathan Timcheck, Alexey Svyatkovskiy, and Neel Sundaresan. 2020. PyMT5: Multi-Mode Translation of Natural Language and Python Code with Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 9052–9065. <https://doi.org/10.18653/v1/2020.emnlp-main.728>
- [11] P. Devanbu. 2015. New Initiative: The Naturalness of Software. In *Proceedings - International Conference on Software Engineering*, Vol. 2. IEEE Computer Society, 543–546. <https://doi.org/10.1109/icse.2015.190>
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL, Vol 1)*. Association for Computational Linguistics, 4171–4186. <https://doi.org/10.18653/v1/n19-1423>
- [13] Angela Fan, Edouard Grave, and Armand Joulin. 2019. Reducing Transformer Depth on Demand with Structured Dropout. <https://doi.org/10.48550/arXiv.1909.11556> arXiv:1909.11556 [cs, stat]
- [14] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Online, 1536–1547. <https://doi.org/10.18653/v1/2020.findings-emnlp.139>
- [15] Michael Fu, Chakkrith Tantithamthavorn, Trung Lu, Van Nguyen, and Dinh Phung. 2022. VulRepair: A T5-based Automated Software Vulnerability Repair. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2022)*. Association for Computing Machinery, New York, NY, USA, 935–947. <https://doi.org/10.1145/3540250.3549098>
- [16] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *International Conference on Learning Representations, ICLR 2021*. Virtual Event, Austria, 1–18. arXiv:2009.08366 [cs]
- [17] Vincent J. Hellendoorn, Christian Bird, Earl T. Barr, and Miltiadis Allamanis. 2018. Deep Learning Type Inference. In *Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM, New York, NY, USA, 152–162. <https://doi.org/10.1145/3236024.3236051>
- [18] José Antonio Hernández López, Martin Weysow, Jesús Sánchez Cuadrado, and Houari Sahraoui. 2023. AST-Probe: Recovering Abstract Syntax Trees from Hidden Representations of Pre-Trained Language Models. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3551349.3556900>
- [19] Jeremy Howard and Sebastian Ruder. 2018. Universal Language Model Fine-tuning for Text Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Melbourne, Australia, 328–339. <https://doi.org/10.18653/v1/p18-1031>
- [20] Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. Learning and Evaluating Contextual Embedding of Source Code. In *Proceedings of the 37th International Conference on Machine Learning*. PMLR, 5110–5121. arXiv:2001.00059 [cs]
- [21] Anjan Karmakar and Romain Robbes. 2021. What Do Pre-Trained Code Models Know about Code?. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 1332–1336. <https://doi.org/10.1109/ase51524.2021.9678927>
- [22] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. <https://doi.org/10.48550/arXiv.1909.11942> arXiv:1909.11942 [cs]
- [23] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Gesheng Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Bulchanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023. StarCoder: May the Source Be with You! <https://doi.org/10.48550/arXiv.2305.06161> arXiv:2305.06161 [cs]
- [24] Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. Linguistic Knowledge and Transferability of Contextual Representations. <https://doi.org/10.48550/arXiv.1903.08855> arXiv:1903.08855 [cs]
- [25] Shangqing Liu, Bozhi Wu, Xiaofei Xie, Guozhu Meng, and Yang Liu. 2023. ContraBERT: Enhancing Code Pre-Trained Models via Contrastive Learning. In *Proceedings of the 45th International Conference on Software Engineering (ICSE '23)*. IEEE Press, Melbourne, Victoria, Australia, 2476–2487. <https://doi.org/10.1109/icse48619.2023.020207>
- [26] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. <https://doi.org/10.48550/arXiv.1907.11692> arXiv:1907.11692 [cs]
- [27] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. 1–16. arXiv:2102.04664 [cs]
- [28] Changan Niu, Chuanyi Li, Bin Luo, and Vincent Ng. 2022. Deep Learning Meets Software Engineering: A Survey on Pre-Trained Models of Source Code. <https://doi.org/10.48550/arXiv.2205.11739> arXiv:2205.11739 [cs]
- [29] Matteo Paltenghi and Michael Pradel. 2021. Thinking Like a Developer? Comparing the Attention of Humans with Neural Models of Code. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 867–879. <https://doi.org/10.1109/ase51524.2021.9678712>
- [30] Cong Pan, Minyan Lu, and Biao Xu. 2021. An Empirical Study on Software Defect Prediction Using CodeBERT Model. *Applied Sciences* 11, 11 (May 2021), 4793. <https://doi.org/10.3390/app11114793>
- [31] David Peer, Sebastian Stabinger, Stefan Engl, and Antonio Rodriguez-Sanchez. 2022. Greedy-Layer Pruning: Speeding up Transformer Models for Natural

- Language Processing. *Pattern Recognition Letters* 157 (May 2022), 76–82. <https://doi.org/10.1016/j.patrec.2022.03.023> arXiv:2105.14839 [cs]
- [32] Matthew Peters, Mark Neumann, Luke Zettlemoyer, and Wen-tau Yih. 2018. Dissecting Contextual Word Embeddings: Architecture and Representation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 1499–1509. <https://doi.org/10.18653/v1/d18-1179>
- [33] Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks. In *Proceedings of the 4th Workshop on Representation Learning for NLP (ReplANLP-2019)*. Association for Computational Linguistics, Florence, Italy, 7–14. <https://doi.org/10.18653/v1/w19-4302>
- [34] Long Phan, Hieu Tran, Daniel Le, Hieu Nguyen, James Annibal, Alec Peltekian, and Yanfang Ye. 2021. CoText: Multi-task Learning with Code-Text Transformer. In *Workshop on Natural Language Processing for Programming (NLP4Prog 2021)*. Association for Computational Linguistics, Online, 40–47. <https://doi.org/10.18653/v1/2021.nlp4prog-1.5>
- [35] Rebecca Russell, Louis Kim, Lei Hamilton, Tomo Lazovich, Jacob Harer, Onur Ozdemir, Paul Ellingwood, and Marc McConley. 2018. Automated Vulnerability Detection in Source Code Using Deep Representation Learning. In *International Conference on Machine Learning and Applications (ICMLA)*. IEEE, Orlando, FL, 757–762. <https://doi.org/10.1109/icmla.2018.00120>
- [36] Hassan Sajjad, Fahim Dalvi, Nadir Durrani, and Preslav Nakov. 2023. On the Effect of Dropping Layers of Pre-trained Transformer Models. *Computer Speech & Language* 77 (Jan. 2023), 101429. <https://doi.org/10.1016/j.csl.2022.101429> arXiv:2004.03844 [cs]
- [37] Rishab Sharma, Fuxiang Chen, Fatemeh Fard, and David Lo. 2022. An Exploratory Study on Code Attention in BERT. In *Proceedings of the 30th IEEE/ACM International Conference on Program Comprehension (ICPC '22)*. Association for Computing Machinery, New York, NY, USA, 437–448. <https://doi.org/10.1145/3524610.3527921>
- [38] Tushar Sharma, Maria Kechagia, Stefanos Georgiou, Rohit Tiwari, and Federica Sarro. 2021. A Survey on Machine Learning Techniques for Source Code Analysis. <https://doi.org/10.48550/arXiv.2110.09610> arXiv:2110.09610 [cs]
- [39] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to Fine-Tune BERT for Text Classification?. In *Chinese Computational Linguistics (Lecture Notes in Computer Science)*, Maosong Sun, Xuanjing Huang, Heng Ji, Zhiyuan Liu, and Yang Liu (Eds.). Springer International Publishing, Cham, 194–206. https://doi.org/10.1007/978-3-030-32381-3_16
- [40] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2020. How to Fine-Tune BERT for Text Classification? <https://doi.org/10.48550/arXiv.1905.05583> arXiv:1905.05583 [cs]
- [41] András Vargha and Harold D. Delaney. 2000. A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics* 25, 2 (June 2000), 101–132. <https://doi.org/10.3102/10769986025002101>
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *International Conference on Neural Information Processing Systems (NeurIPS)*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5998–6008. arXiv:1706.03762 [cs]
- [43] Yue Wang, Weishi Wang, Shafiq Joty, and Steven C.H. Hoi. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 8696–8708. <https://doi.org/10.18653/v1/2021.emnlp-main.685>
- [44] Hongwei Wei, Guanjun Lin, Lin Li, and Heming Jia. 2021. A Context-Aware Neural Embedding for Function-Level Vulnerability Detection. *Algorithms* 14, 11 (Nov. 2021), 335. <https://doi.org/10.3390/a14110335>
- [45] Frank Wilcoxon. 1992. Individual Comparisons by Ranking Methods. In *Breakthroughs in Statistics: Methodology and Distribution*, Samuel Kotz and Norman L. Johnson (Eds.). Springer, New York, NY, 196–202. https://doi.org/10.1007/978-1-4612-4380-9_16
- [46] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2020. XLNet: Generalized Autoregressive Pretraining for Language Understanding. <https://doi.org/10.48550/arXiv.1906.08237> arXiv:1906.08237 [cs]
- [47] Michihiro Yasunaga and Percy Liang. 2021. Break-It-Fix-It: Unsupervised Learning for Program Repair. In *International Conference on Machine Learning*. PMLR, 12. arXiv:2106.06600 [cs]
- [48] He Ye, Matias Martinez, and Martin Monperrus. 2022. Neural Program Repair with Execution-Based Backpropagation. In *Proceedings of the 44th International Conference on Software Engineering (ICSE '22)*. Association for Computing Machinery, New York, NY, USA, 1506–1518. <https://doi.org/10.1145/3510003.3510222>
- [49] Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q. Weinberger, and Yoav Artzi. 2021. Revisiting Few-sample BERT Fine-tuning. In *NeurIPS 2021*. 1–22. arXiv:2006.05987 [cs]
- [50] Gang Zhao and Jeff Huang. 2018. DeepSim: Deep Learning Code Functional Similarity. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, Lake Buena Vista FL USA, 141–151. <https://doi.org/10.1145/3236024.3236068>
- [51] Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. 2019. Devign: Effective Vulnerability Identification by Learning Comprehensive Program Semantics via Graph Neural Networks. In *International Conference on Neural Information Processing Systems (NeurIPS)*. Curran Associates, Inc., Vancouver, Canada., 11. arXiv:1909.03496 [cs]

Received 2023-02-02; accepted 2023-07-27