

Assemble Foundation Models for Automatic Code Summarization

Jian Gu

University of Zurich
Zurich, Switzerland
e-mail: gu@ifi.uzh.ch

Pasquale Salza

University of Zurich
Zurich, Switzerland
e-mail: salza@ifi.uzh.ch

Harald C. Gall

University of Zurich
Zurich, Switzerland
e-mail: gall@ifi.uzh.ch

Abstract—Automatic code summarization is beneficial to software development and maintenance since it reduces the burden of manual tasks. Currently, artificial intelligence is undergoing a paradigm shift. The foundation models pretrained on massive data and finetuned to downstream tasks surpass specially customized models. This trend inspired us to consider reusing foundation models instead of learning from scratch. Based on this, we propose a flexible and robust approach for automatic code summarization based on neural networks. We assemble available foundation models, such as CodeBERT and GPT-2, into a single model named AdaMo. Moreover, we utilize Gaussian noise as the simulation of contextual information to optimize the latent representation. Furthermore, we introduce two adaptive schemes from the perspective of knowledge transfer, namely continuous pretraining and intermediate finetuning, and design intermediate stage tasks for general sequence-to-sequence learning. Finally, we evaluate AdaMo against a benchmark dataset for code summarization, by comparing it with state-of-the-art models.

Index Terms—transfer learning, adaptive scheme, Transformer, Gaussian noise, code summarization

I. INTRODUCTION

In the process of software development and maintenance, decent comments are crucial to program comprehension and could reduce the burden of directly interpreting the code [1]. By reading and writing code comments, software developers instantly load and save the working context, such as requirements and implementation details. However, it is laborious to manually write comments in a consistent and decent style and keep them synchronized with the code simultaneously. It seems inevitable to study the automatic way of generating high-quality code comments, namely “code summarization”, which could save massive effort and time.

Code summarization is the task of generating natural language descriptions for the given code snippets. In practice, it serves numerous daily activities in software development and maintenance, such as recording implementation details, templating package descriptions, and describing code changes for version updates [2]. Its automatic solutions could improve productivity in software activities effectively. There are several solutions proposed for automatic code summarization, and among them, neural models are the most valued for their unique generative ability. These machine learning solutions significantly rely on the knowledge learned from a large code corpus. Correspondingly, neural models require massive data and intensive computational resources.

In this paper, we build a transfer learning model **ADAMO** composed of relevant foundation models for code summarization and also introduce Gaussian noise to enhance the latent representations. Our experimental results showed that ADAMO defeated state-of-the-art (SOTA) models.

Recent work tend to introduce structure information extracted from the parsed results of code data, e.g., Abstract Syntax Tree (AST) and Control Flow Graph (CFG), as a modality complement. The representatives are SOTA models, namely BASTS [3] and SiT [4]. Instead of lightly processing the code data as a token sequence, they parse it into tree or graph forms, thereby taking heavy computation consumption. Compared with SOTA models, our approach is merely based on the token sequence. As token-based neural models proved to be effective in handling sequential data, especially in Natural Language Processing (NLP), we want to leverage on them and see how far we can go. This brings us an advantage over baseline models that probably requires fully parsable code, whereas our model works as well on corrupted or partial code.

To make up for the disadvantages that our model only learns from less information, our model emphasizes reusing pretrained checkpoints for better parameter initialization [5]. With the popularity of transfer learning, foundation models are also emerging [6]. They are generally trained on intensive data at scale and then adapted to various target tasks. Based on that, we leverage pretrained models instead of training a model from scratch. We assemble the representative encoder model and decoder model, e.g., BERT [7] and GPT [8], and then train their assembly. Furthermore, considering that adaptive schemes are already proposed in transfer learning, as the complement to the standard “*pretraining* and then *finetuning*” paradigm, we further design intermediate tasks to adopt continuous pretraining and intermediate finetuning on code summarization.

To summarize, the main contributions of this paper are:

- proposing an effective and straightforward approach for code summarization, by assembling foundation models;
- adopting the Gaussian noise emitter as a simulation of contextual information for better latent representations;
- introducing adaptive transfer learning schemes as options to further raise the upper bound of model performance;
- designing intuitive intermediate stage tasks for code summarization and general sequence-to-sequence learning.

The rest of this paper is structured as follows. In Section II, we introduce the background knowledge. In Section III, we present our transfer learning model and adaptive schemes. Section IV describes the experimental setup, whereas Section V presents and discusses the results. Section VI surveys the related work, and Section VII concludes with a summary of the findings and contributions and an outlook on future work.

II. BACKGROUND

In this section, we describe the background of our work, such as Transformer models, e.g., BERT [7] and GPT [8], transfer learning and adaptive schemes.

A. Transformer Models

The standard Transformer is an encoder-decoder neural model that merely relies on the attention mechanism [9] as the main component. We call other models inspired by this simple and effective way of building networks as Transformers.

The architecture of Transformer is a stack of six encoder layers mingled with a stack of six decoder layers. In the encoder layer, there are a self-attention sublayer and a feed-forward network. Instead, in the decoder layer, there is an extra encoder-decoder attention sublayer. The self-attention sublayer is for the connections within encoder layers or decoder layers, while the encoder-decoder attention sublayer is for the connections between the last encoder layer and each decoder layer. Both the self-attention sublayer and encoder-decoder attention sublayer adopt multiple attention units to quantify the importance of each part of the input data in the same manner.

Due to the superiority of Transformers in the computation complexity and flexibility over prior models, they are popular, especially in the field of NLP. Besides, its capability of computation parallelization allows more intense use of massive data, which led to the development of large-scale pretrained models, namely foundation models [6]. In this paper, we consider two examples of Transformers in particular: BERT [7] for discriminative tasks, and GPT [8] for generative tasks.

Bidirectional Encoder Representations from Transformers (BERT). BERT [7] is the typical encoder model inspired by Transformer and follows the idea of merely using the attention mechanism to build the neural model.

In terms of model architecture, BERT stacks 12 encoder layers (24 layers for its large version), but no decoder layers. BERT-like models are used for *discriminative tasks*. Each encoder layer of BERT consists of a self-attention sublayer and a feed-forward network, the same as Transformer. In the self-attention sublayer, each token can attend context to its left and right, so the attention is referred to as “bidirectional”.

Masked Language Modeling (MLM) is a common objective to train BERT-like models. In MLM, the model is trained to predict 15% randomly masked tokens based on the entire context, namely the tokens that occurred on both two sides. The MLM objective works in pair with the bidirectional self-attention, to guide the model to learn a decent representation of the input. RoBERTa [10], as its optimized version, achieves

further improvements by dynamically masking tokens and offering a bigger data capacity. CODEBERT [11] is a specialized RoBERTa for the representation of code data.

Generative Pretrained Transformer (GPT). GPT [8] is the typical decoder model inspired by Transformer that mainly uses the attention mechanism to build the neural model.

In terms of model architecture, GPT only stacks 12 decoder layers (24, 36, or 48 layers for its medium, large, and extra large versions), but no encoder layers. GPT-like models are used for *generative tasks*. Each decoder layer in GPT consists of a masked self-attention sublayer and a feed-forward network, which is different from Transformer. Besides, there is no encoder-decoder attention sublayer. In the masked self-attention sublayer, each token only attends context to its left. Thus the attention is referred to as “constrained”.

The only self-supervised training objective of GPT is Casual Language Modeling (CLM). In CLM, the model learns to predict 15% randomly masked tokens when given the partial context, namely the tokens that occurred only on their left sides. The CLM objective works in pair with the constrained self-attention, to guide the model to generate a proper sentence as the output. Its enhanced versions, i.e., GPT-2 [12], show stronger capability in text generation.

B. Transfer Learning and Adaptive Schemes

Different from general learning modes, *Transfer Learning* is the learning paradigm of transferring the knowledge learned from other data to the new data for given domains or tasks [13]. It relaxes the requirements on data amount, data distribution, and computation capability [14]. Besides the conventional “*pretraining* and then *finetuning*” practice, adaptive schemes are applicable to transfer learning models.

The intuition behind transfer learning is that the beforehand self-supervised training for data representation promises general and reliable initialization, i.e., knowledge of data distribution, and then benefits the performance of specific downstream tasks, such as raising the upper bound or accelerating convergence. Compared with starting from scratch, transfer learning could improve the sample efficiency since it reduces the consumption required on data and computational resources [15].

Based on transfer learning, models for various tasks could origin from the same pretrained language model. As an extended definition of pretrained language model, the model that is “*trained on broad data at scale and can be adapted to a wide range of downstream tasks*” is named “*foundation model*” [6]. In NLP, representative examples include BERT and GPT, which are trained on large corpora of text and then adapted to a wide range of downstream tasks, e.g., machine translation, question answering, and sentiment analysis.

The standard methodology of transfer learning is *pretraining* a model on a large corpus of unlabeled data and then *finetuning* it on a small supervised dataset. In the pretraining stage, the model is usually trained in a *self-supervised learning* manner, where the unlabeled data is sufficient for the objective, therefore the pretraining data is usually extensive and readily available. Instead, during the finetuning stage, the model is trained in a

supervised learning manner where the ground truth is required. The data used for the target (downstream) task is supervised data, and its quality matters the most, so the data amount is usually limited and the cost is more expensive. Even though it has been conventional to pretrain and then finetune models, there are still strategies proposed to further improve the data adaptability to target domains or tasks, which mainly work between the usual pretraining and finetuning stages.

Continuous Pretraining (CP) is defined as tailoring a model to another data domain, or even the data of a target task, via the second phase of pretraining [16]. It has been shown that domain-adaptive pretraining, namely adapting the model to the data of the same domains, improves the performance. Similarly, adapting the model for designated tasks, called task-adaptive pretraining, leads to performance gain as well. Moreover, multiphase adaptive pretraining, e.g., domain-adaptive training followed by the task-adaptive one, promises a larger gain.

Intermediate Finetuning (IF) benefits a pretrained model by introducing intermediate tasks during the additional training stage, as the warmup activities before training for the target task [17]. However, the characteristics of intermediate tasks could affect the effectiveness of adaptive finetuning. Divided by difficulty, intermediate tasks could be either simple or complex. The simple intermediate tasks are close to learning the low-level skills such as preserving the raw content and detecting the shallow attributes, such as verb tenses or sentences length. They would only affect the model performance slightly. In contrast, complex intermediate tasks are generally rather beneficial to promote the model, e.g., natural language inference [18], and question answering [19]. Thus, they expect the model to have strong capabilities such as perceiving interrelations.

III. APPROACH

This section describes the model architecture and its components and how we applied transfer learning. Then, we introduce adaptive schemes [20], e.g., continuous pretraining, and intermediate finetuning.

A. Model Architecture

Our transfer learning model follows the encoder-decoder architecture for general sequence-to-sequence learning [21], which has stacked encoding layers as the encoder and stacked decoding layers as the decoder.

Instead of training from scratch, building a model by adopting existing foundation models supports reusing the pretrained weights for model initialization [22]. Thereby, we propose such a model that adopts CODEBERT [11], e.g., a specialized RoBERTa for code representation, as the encoder, and GTP-2 [12], e.g., an enhanced GPT for text generation, as the decoder. In addition, we place an additive white Gaussian noise (AWGN) emitter between the encoder and the decoder as an option to further optimize the latent representation. We name such a transfer learning model **ADAMO**, standing for **Adaptive Model with reference to the adaptive schemes it supports**.

Considering the input data is code snippets, and output data is code comments, our model adopts CODEBERT as the




Figure 1. The ADAMO architecture.




Figure 2. The Adaptive schemes.

encoder and GTP-2 as the decoder. CODEBERT can capture the general representations of code data, and similarly, GTP-2 performs well in English text generation. We make no changes to the metadata of the encoder model and decoder, such as model structure and hyperparameters. Besides, the noise emitter between encoder and decoder is expected to generate Gaussian noise as the rough simulation of contextual information of the input data [23]. As shown in Fig. 1, the red blocks are input data and output data, and the dashed block is the architecture of our model, whose components are represented as blue blocks.

We define the processor manipulating the latent representations between encoder and decoder as *Refiner*. Furthermore, we define the latent representations of encoder and decoder as *Parallel Representations*. In our model architecture, encoder and decoder both reuse the network design of existing language models and are initialized with the pretrained weights. Therefore the parallel representations of both sides are not naturally coordinated with each other. The refiner is thereby introduced to optimize the latent representations and reduce their incoordination.

Prior work in natural language translation found it could be equivalent to introducing multiple encoders to capture contextual information when accumulating Gaussian noise to the latent representations [23]. Gaussian noise is a basic noise model, following the normal distribution, commonly used to mimic the effect of random processes in nature. Considering both code and text data are sequential, and the context information matters as well in code data, we adopt the AWGN emitter to simulate the context-aware setting of multiple encoders.

Our ADAMO model is mainly composed of a Transformer encoder and decoder, and its architecture naturally supports practices applicable in transfer learning, such as the adaptive schemes. As shown in Fig. 2, we introduce adaptive schemes, e.g., continuous pretraining and intermediate finetuning, between the common pretraining and finetuning stages, to make the model more adaptive to the data of specific domains or tasks. The adaptive schemes are expected to bring improvements by raising the upper bound of model performance. They do not necessarily appear simultaneously in our approach, but their combination is feasible.

B. Adaptive Schemes

Based on the standard transfer learning paradigm of “*pretraining* and then *finetuning*”, we introduce adaptive schemes

for better data adaptability to target domains or tasks, e.g., continuous pretraining and intermediate finetuning.

For continuous pretraining, irrespective whether it is Domain-Adaptive (DA) or Task-Adaptive (TA), there exist no obvious differences with the general pretraining stage. However, DA pretraining specifies that the unlabeled data to be used must come from a related domain, while TA pretraining directly utilizes the unlabeled data of the given task. For instance, assuming that we would use a dataset of Physics papers to finetune a pretrained language model, depending on the concrete data to use, continuous pretraining is identified as domain-adaptive or task-adaptive. If the unlabeled data is not from the given dataset but another related dataset, such as the dataset of general scientific papers, it is DA pretraining. On the contrary, if the unlabeled data is from the given dataset, it is referred to as TA pretraining.

For both DA and TA, we set Masked Language Modeling (MLM) as the objective of the encoder, and Casual Language Modeling (CLM) as the objective of the decoder. Both MLM and CLM involve randomly masking parts of the sequential data, then training the model to predict the missing tokens correctly. For the prediction, MLM allows the model to consider the entire context, i.e., the data occurring to both left and right sides. Instead, CLM only allows the model to consider the partial context, i.e., the data occurring to the left side only. Specifically, for DA pretraining, we pretrain the model with massive unlabeled data from another related dataset, separately train only either encoder or decoder, or train them both. For TA pretraining, we pretrain the model almost in the same way but directly using the unlabeled data from the given dataset.

Domain-Adaptive: The adaptation to the distribution of the related domain of the target data in the model migration.

Task-Adaptive: The adaptation to the distribution of the target data when the model migrates from the source data.

To the best of our knowledge, there exist no ready-made intermediate tasks designed for code summarization yet, and none for general sequence-to-sequence learning. Inspired by text editing [24], we propose that the summarization task could be seen as the fusion of two novel stage tasks: one is to insert tokens that occurred in the target text but not in the source text; the other one is to delete tokens that occurred in the source text but not in the target text. Also, the stage tasks concern the order of tokens as well because reordering tokens is a mandatory implicit operation. Considering this, we regard the changes in both the occurrence and order of tokens in the summarization task as concept shifts. We name the first stage task as Concept Extrapolation (CE) and the second one as Concept Interpolation (CI). In addition, one intuitive but straightforward idea is to directly annotate the tokens in the target text with different marks based on whether they are available in the source text. We propose this idea as the comparison task and name it as Concept Annotation (CA). Not limited to *code* summarization,

CA, CE, and CI are applicable to *text* summarization, even general sequence-to-sequence learning.

For a given sentence, both CE and CI cover the operation of reordering tokens, however, CE is more like a typical complex task while CI seems relatively simple. The reason is that the model has to associate and append tokens that never appeared in the sentence in the case of CE. Instead, for CI the model merely determines whether each existing token in the sentence should be kept or dropped. Concept Annotation (CA) is a bit different from others because it only annotates tokens and introduces no information about the order of the target text. We regard it as the most straightforward intermediate task. The number of intermediate tasks is suggested to be 1 because when the amount of involved tasks is less than a critical point, e.g., 15, then the fewer tasks, the better the performance [25].

Concept Annotation: The stage task masks the tokens in the target sequence with different tokens based on whether they have ever appeared in the source sequence or not.

Concept Extrapolation: The stage task masks the tokens in the target sequence that have already appeared in the source sequence, to drive the model to connect new tokens.

Concept Interpolation: The stage task masks the tokens in the target sequence that have never appeared in the source sequence, to drive the model to discard old tokens.

IV. EXPERIMENTS

To check the effectiveness of our approach which takes reusing pretrained language models as a prerequisite, we first experiment with the transfer learning assembly itself and then introduce adaptive schemes, e.g., continuous pretraining and intermediate finetuning. In the context of the study, we thereby formulate the research questions as follows.

- RQ1** *How well does our transfer learning model perform in code summarization?*
- RQ2** *How well does code summarization benefit from the AWGN refiner?*
- RQ3** *How well does code summarization benefit from adaptive continuous pretraining?*
- RQ4** *How well does code summarization benefit from adaptive intermediate finetuning?*

The replication repository is published online ¹, including implementations, configuration, results, and the guidance on reproducing experiments. The implementations are in PYTHON, using PYTORCH [26] and TRANSFORMERS [27]. The scoring function invokes the third-party evaluation package [28]. The experiments are conducted on a machine with an AMD EPYC 7702 CPU with 16 GB RAM, and a single Nvidia Tesla V100 GPU with 32 GB memory.

In the following, we describe the methodology applied to answer the above-mentioned research questions.

¹<https://github.com/jianguda/afm4acs>

A. Baselines

We compare our model with modern SOTA models, e.g., BASTS [3] and SiT [4]. They both utilize structure information to improve the code summarization and outperformed other prior models but have never been compared to each other.

BASTS parses the code into an AST and converts it into smaller trees, capturing the local non-linear syntax information from them. It splits the code snippet based on its blocks in the dominator tree of the CFG and generates a split AST for each code split. Each split AST will be encoded in the tree-form separately for concatenation with the usual code embedding.

SiT introduces a structure-induced self-attention mechanism to capture the information on syntax structure, data flow as well as data dependency from AST. It parses the code snippet into an AST, then transforms the tree into adjacency matrices of three views, and finally combines them evenly to complement the information of sequential code data.

A crucial difference between our transfer learning model and baseline models is that the baseline models are not compatible with adaptive schemes. It is caused by the extensibility of our model. We successively use model tags BASTS and SiT to indicate the introduced baseline models.

B. Corpus

In our experiments, we use three datasets in total. All datasets consist of numerous pairs of code snippets and code comments in JAVA or PYTHON. All the experiments are conducted on JAVA and PYTHON separately.

In code summarization, the commonly used evaluation is the one used in SiT [4], in which JAVA and PYTHON are selected as the representative research targets considering their language popularity and data richness. We thereby name this benchmark as SiT dataset, and we conduct experiments on that dataset as well. The other baseline model, i.e., BASTS [3], adopted a different benchmark. For the sake of clarity, we name that benchmark as BASTS dataset. Considering the difficulty of reproducing BASTS, due to its weak adaptability of the preprocessing pipeline, especially for the performance loss that happens in switching the benchmark, we evaluate our model on the BASTS dataset for fair comparisons.

In addition to BASTS and SiT datasets, we introduce the CodeSearchNet dataset [29], shortly referred to as CSN, to support the experiments for adaptive schemes. The CSN dataset is collected from publicly available open-source repositories in GITHUB. The full dataset contains over two million snippet-comment pairs spanning multiple programming languages. To keep consistent with other datasets, we only utilize its JAVA and PYTHON data. We follow the given data partition of the CSN dataset but merge the given test set into the given training set, considering we would never use the given test set.

Actually, the BASTS dataset originates from earlier data sources [29], [30], so does the SiT dataset [31], [32]. Even though the PYTHON data in the BASTS dataset are taken from the CSN dataset, these two datasets are never used for the same model, so there is no potential data leak. Meanwhile, all the data in the SiT dataset originates from GITHUB, so

TABLE I
STATISTICS OF THE SNIPPET-COMMENT PAIRS IN THE CORPUS

Dataset	BASTS		SiT		CSN	
	JAVA	PYTHON	JAVA	PYTHON	JAVA	PYTHON
Train	415,395	216,436	69,708	57,203	170,106	265,734
Valid	12,885	12,119	8,714	19,067	10,955	14,918
Test	13,237	12,767	8,714	19,066	-	-
Total	441,517	241,322	87,136	95,336	181,061	280,652

there might be a data overlap between SiT and CSN datasets. However, in our experimental design, we never directly use data pairs in CSN for sequence-to-sequence learning, so no potential data leak exists as well.

As shown in Table I, the first column presents datasets, and the second to ninth columns the statistics on the pairs of code snippets and code comments, per language and dataset.

C. Evaluation Metrics

Considering BLEU [33], METEOR [34], and ROUGE [35] are the most common evaluation metrics in code summarization, we thus use them to evaluate our experiments [36].

The definitions of BLEU, METEOR, and ROUGE are based on the same scenario, i.e., for candidate sentences in a corpus, there are a set of reference sentences that corresponds to each of them. In their implementations, the penalty factor ρ is for rational adjustments. Besides, their computations rely on the Precision score and the Recall score, short as P and R :

$$P_n = \frac{\#\text{gram}_n(c, r)}{\#\text{gram}_n(c)}, R_n = \frac{\#\text{gram}_n(c, r)}{\#\text{gram}_n(r)}$$

where c and r are the candidate sentence and reference sentence, respectively, and $\#\text{gram}$ is the number of overlapped n -grams.

BLEU. The BLEU score computes the averaged percentage of n -gram matches between the candidate sentence and the reference sentence, typically unigrams through 4-grams:

$$\text{BLEU} = \rho \cdot \exp\left(\frac{1}{N} \sum_{n=1}^N \log P_n\right) \quad (1)$$

There are two implementation versions of the BLEU score. In one case, as proposed initially, the computation is on the corpus level, namely computing one score for the whole corpus, the score is ‘‘corpus BLEU’’, referred to as C-BLEU. In the other case, the computation is on the sentence level, i.e., computing one score for each sentence and taking their arithmetic mean as the final score, the score is ‘‘sentence BLEU’’, referred to as S-BLEU.

METEOR. The METEOR score builds alignments on unigrams between the candidate sentence and the reference sentence, but during the process, alignments are prioritized on the longer n -grams:

$$\text{METEOR} = \rho \cdot \frac{10P_n R_n}{R_n + 9P_n} \quad (2)$$

Roughly speaking, the penalty factor depends on the actual alignment. When the alignment is mainly on the unigrams, the penalty could be rather heavy because of the low similarity.

ROUGE. The ROUGE score counts the overlaps of n -grams or Longest Common Sequences (LCS) between the candidate sentence and the corresponding reference sentence. Taking the n -grams as an example, the formula is as follows:

$$\text{ROUGE} = \frac{2P_n R_n}{R_n + P_n} \quad (3)$$

The most common ROUGE scores are ROUGE-1, ROUGE-2, and ROUGE-L. Their distinction lies on the overlapped target to be counted, e.g., unigrams, 2-grams, or LCS.

In our experiments, we compute C-BLEU, S-BLEU, METEOR, and ROUGE-L to quantify the quality of the summarization results. These metrics are in the range of $[0, 1]$ and will be reported in the percentage form. The larger the value, the better the effect. Based on the conclusion of the empirical study of automatic evaluation metrics for code summarization, we give priority to METEOR [37].

D. Methodology for RQ1

There have been numerous pretrained language models proposed for general code-related tasks [38], but there is not yet prior work on leveraging existing checkpoints for code summarization. Meanwhile, what the performance margin is between such transfer learning models and existing well-designed models is still unknown to the community.

First, we design experiments to compare ADAMO and SOTA models on the same datasets. Considering both the BASTS and SiT models involve complicated preprocessing steps to extract structure information, it is more fair and intuitive to experiment with ADAMO on corresponding datasets directly. Moreover, the potential performance reduction caused by scripts migration may not be easily avoided. In detail, we experiment with ADAMO on the BASTS dataset to compare with BASTS, and on the SiT dataset to compare with SiT.

We conducted two experiments to build an intuitive understanding of model performance. In one group of experiments, we directly evaluate the model without the finetuning step. In this way, we could estimate how well the pretrained model performs in itself, i.e., without any knowledge of the downstream task, namely *zeroshot learning* [39]. In the other group, we evaluate the encoder-decoder model with the finetuning step. Here we assign a roughly equivalent time budget, e.g., 24 hours, to BASTS, SiT, and ADAMO for a fair comparison.

The tag used for zeroshot learning model is `AdaMo-0shot`, and for the normal finetuning is `AdaMo-basic`. In our experiments, the noise emitter is temporarily turned off.

E. Methodology for RQ2

To study the effectiveness of the AWGN emitter and how its configuration affects the performance of our transfer learning model, we enable the Gaussian noise emitter but with different configurations. Besides, we run each experiment for the same epoch number with the noise emitter off.

The experimental design is almost the same as for RQ1 because we merely adjust the variance to specific values,

without making any other changes. The experiments are on the SiT dataset, which was commonly used in prior work.

The tag used for the noise model is `AdaMo-noise`. We use labels in the form of `AdaMo-noise[σ]` to distinguish different settings, where the standard deviation σ could be 0.1, 0.2, 0.3, 0.4 or 0.5. For example, the label for the AWGN emitter, when the standard deviation σ is 0.1, is `AdaMo-noise[0.1]`.

F. Methodology for RQ3

Results from NLP indicate that the second phase of pretraining in the domain leads to performance gains [16]. Therefore, it has a promising potential improvement for code summarization.

As introduced in Section III-B, we adopt Domain-Adaptive (DA) and Task-Adaptive (TA) pretraining as the additional procedure of tailoring the pretrained model to the data of target domains or tasks. For DA pretraining, we utilize the CSN dataset to train encoder, decoder, or them both separately. For TA pretraining, we separately train encoder, decoder, or both on the SiT dataset. The BASTS dataset is excluded since it has a times amount of JAVA data to CSN and their PYTHON data show an extreme overlap.

As the experimental settings of continuous pretraining, we train the model in both the DA and TA way. Meanwhile, we separately train the encoder, the decoder, and them both in all experiments. To check the performance of the adaptive model after the adaptive continuous pretraining scheme, we offer 24 hours for the additional training phase. That is, the encoder, the decoder, or them both are pretrained for 24 hours first, and then ADAMO itself is trained for 24 hours. Meanwhile, to validate whether continuous pretraining usually requires a longer time like the normal pretraining, we offer a richer time budget, e.g., 48 hours, as the reference.

The tag used for continuous pretraining is `AdaMo-CP`. We use labels in the form of `AdaMo-CP[ADAPTIVE][OBJ]` to distinguish different settings, where `ADAPTIVE` can be DA or TA, and `OBJ` could be `m1m`, `c1m` or `both`. For example, the label of continuous pretraining only the encoder in the domain-adaptive manner is referred to as `AdaMo-CP[DA][m1m]`.

G. Methodology for RQ4

For NLP tasks that require high-level inference and reasoning capabilities, it proves to be effective to introduce relevant tasks with massive data as the intermediate finetuning [17].

As introduced in Section III-B, we want to check whether the intermediate finetuning might benefit our transfer learning approach. Meanwhile, we expect the ideas of domain and task adaptations could bring potential performance improvements.

As the experimental settings of intermediate finetuning, we conduct two sets of experiments where the model is trained in the domain-adaptive or task-adaptive ways. In each set, there are three experiments in which CA, CE or CI is adopted as the intermediate stage task. We train the model for the stage task on the CSN dataset as the domain-adaptive setting, and on the SiT dataset as the task-adaptive setting. The BASTS dataset is excluded since it has a times amount of JAVA data to CSN and their PYTHON data show an extreme overlap.

TABLE II
COMPARATIVE RESULTS ON CODE SUMMARIZATION

Dataset	Model	JAVA				PYTHON			
		C-BLEU	S-BLEU	METEOR	ROUGE	C-BLEU	S-BLEU	METEOR	ROUGE
BASTS	BASTS	27.82%	34.22%	22.86%	45.78%	2.33%	14.18%	8.65%	20.87%
	AdaMo-0shot	0.00%	1.80%	0.10%	0.24%	0.00%	1.78%	0.15%	0.36%
	AdaMo-basic	31.38%	37.64%	25.50%	49.00%	5.19%	16.46%	12.51%	27.31%
SiT	SiT	38.01%	44.96%	27.09%	53.54%	26.46%	33.81%	21.37%	41.18%
	AdaMo-0shot	0.00%	1.79%	0.12%	0.22%	0.00%	1.89%	0.07%	0.13%
	AdaMo-basic	40.49%	45.30%	28.19%	53.99%	26.52%	33.85%	21.68%	41.25%

To investigate the performance of ADAMO after the intermediate finetuning scheme, we offer 24 hours for the additional training phase. For the domain-adaptive and task-adaptive settings, they both have additional 24 hours for intermediate finetuning, apart from the 24 hours for the normal finetuning.

The tag used for intermediate finetuning is AdaMo-IF. We use labels in the form of AdaMo-IF[ADAPTIVE][PROXY] to distinguish different settings, where ADAPTIVE could be DA or TA, and PROXY could be CA, CE or CI. For example, the label of the CA stage task for domain-adaptive is referred to as AdaMo-IF[DA][CA].

V. RESULTS

In this section, we present the results of our experiments to answer the research questions. In the following comparisons, the most competitive results are highlighted in bold.

A. Results of RQ1

The results for RQ1 are the summarization scores of our basic transfer learning approach, in comparison with the baseline models introduced above.

As shown in Table II, on both BASTS and SiT datasets, AdaMo-0shot reaches the lowest scores. On the contrary, AdaMo-basic always performs better than baseline models on all metrics. On the one hand, it shows the incoordination of parallel representations damages the zeroshot ability of AdaMo. On the other hand, it proves the strategy of directly reusing the well-behaved model structures and their trained weights is effective, meanwhile, the training work of AdaMo-basic is not merely for tuning the encoder and decoder, but also to overcome the incoordination issue of parallel representations.

Considering that BASTS and SiT are separately reproduced in their respective dataset, it is infeasible to compare their performance directly. With the reference to AdaMo-basic, it outperforms BASTS significantly but only obtains moderate advantages over SiT, therefore the performance of SiT is closer to ADAMO than BASTS, in terms of the model effects.

RQ1 – Takeaway: Our transfer learning approach beats the SOTA models smoothly when given the same time budget. Even though ADAMO reuses the existing language models and their trained weights, it is indispensable to train their assembly on the data of target domain or for the target task.

B. Results of RQ2

The results for RQ2 are the summarization scores of ADAMO accumulating varied intensities of Gaussian noise, compared with our basic transfer learning approach.

TABLE III
RESULTS OF THE MODEL WITH DIFFERENT AWGN REFINERS

Model	JAVA				PYTHON			
	C-BLEU	S-BLEU	METEOR	ROUGE	C-BLEU	S-BLEU	METEOR	ROUGE
AdaMo-basic	40.49%	45.30%	28.19%	53.99%	26.52%	33.85%	21.68%	41.25%
AdaMo-noise[0.1]	39.04%	43.41%	26.86%	51.67%	25.21%	32.55%	20.79%	39.45%
AdaMo-noise[0.2]	40.82%	45.33%	28.18%	53.86%	26.65%	34.03%	24.95%	41.88%
AdaMo-noise[0.3]	40.52%	45.35%	28.25%	54.06%	26.80%	34.05%	21.92%	41.67%
AdaMo-noise[0.4]	40.37%	44.92%	27.93%	53.36%	26.78%	34.04%	21.89%	41.71%
AdaMo-noise[0.5]	40.56%	45.00%	28.04%	53.49%	26.74%	34.05%	21.91%	41.64%

As shown in Table III, the effectiveness of the AWGN emitter relies on whether or not the noise emitter is well configured. We observed that with increasing values of σ until it is 0.5, the effects decrease first and then increase. When the standard deviation σ is 0.3, ADAMO obtains the best performance. After that, the results start to fluctuate and enter the downtrend on JAVA or a stable state on PYTHON.

Comparing the effects of Gaussian noise on specific languages, the results of AdaMo-basic improve on all evaluation metrics only when the σ is set to 0.3 on JAVA. On the contrary, results of almost all versions of AdaMo-noise are always better than AdaMo-basic on PYTHON. The cause is believed to lie in the differences in language expressiveness. The code snippets in PYTHON are more intuitive than those in JAVA, and also the PYTHON programs are more close to the common English expressions. English is a natural language while programming languages are artificial or constructed languages. Considering that AWGN is commonly used to mimic the random processes in nature and that PYTHON programs are more close to the English text than JAVA in terms of coding rules, the Gaussian noise should be more effective when the data itself is more natural, or less artificial.

Overall, the improvements brought by the Gaussian noise are reliable, but it is not conclusive enough. One reason could be the distinctions in data characteristics. The Gaussian noise could not perfectly simulate the context information of the code data, as it does in the text data. The code data is more artificial while the text data is more natural, therefore the effectiveness of Gaussian noise is reduced. The other reason could be the natural difference between the machine translation task and the code summarization task. Code summarization is a hybrid task of text summarization and machine translation because the code snippets are strictly written in compiler-oriented grammar rules, but the code comments flexibly follow the natural grammar rules, and meanwhile, the information that existed in the source data is largely reduced after summarization.

RQ2 – Takeaway: The Gaussian noise benefits the model when the noise emitter is well configured. The effects are more accessible if the data itself is closer to the natural language. However, the improvements of AWGN are reliable for code summarization but not very obvious.

C. Results of RQ3

The results for RQ3 show the summarization scores of ADAMO when applying the continuous pretraining scheme, by comparing with our basic transfer learning approach. As

TABLE IV
RESULTS OF THE CONTINUALLY PRETRAINING SCHEME (24 HOURS)

Model	JAVA				PYTHON			
	C-BLEU	S-BLEU	METEOR	ROUGE	C-BLEU	S-BLEU	METEOR	ROUGE
AdaMo-basic	40.49%	45.30%	28.19%	53.99%	26.52%	33.85%	21.68%	41.25%
AdaMo-CP [DA] [mlm]	40.50%	45.05%	27.98%	53.40%	26.47%	33.83%	21.55%	40.99%
AdaMo-CP [DA] [c1m]	40.18%	45.22%	28.20%	53.94%	26.35%	33.68%	21.73%	41.41%
AdaMo-CP [DA] [both]	40.62%	45.05%	28.16%	53.82%	26.35%	33.61%	21.46%	40.98%
AdaMo-CP [TA] [mlm]	40.61%	45.27%	28.08%	53.72%	27.10%	34.43%	22.25%	42.53%
AdaMo-CP [TA] [c1m]	40.01%	44.97%	28.19%	53.96%	26.53%	33.80%	21.89%	41.66%
AdaMo-CP [TA] [both]	40.45%	44.93%	27.98%	53.64%	26.86%	34.21%	22.16%	42.37%

TABLE V
RESULTS OF THE CONTINUALLY PRETRAINING SCHEME (48 HOURS)

Model	JAVA				PYTHON			
	C-BLEU	S-BLEU	METEOR	ROUGE	C-BLEU	S-BLEU	METEOR	ROUGE
AdaMo-basic	40.49%	45.30%	28.19%	53.99%	26.52%	33.85%	21.68%	41.25%
AdaMo-CP [DA] [mlm]	40.61%	45.18%	28.03%	53.49%	26.28%*	33.56%*	21.46%*	40.85%*
AdaMo-CP [DA] [c1m]	39.82%*	44.97%*	28.10%*	53.95%	26.34%*	33.67%*	21.81%	41.69%
AdaMo-CP [DA] [both]	40.39%*	44.94%*	27.97%*	53.40%*	26.50%	33.76%	21.78%	41.53%
AdaMo-CP [TA] [mlm]	40.55%*	45.37%	28.15%	53.97%	26.96%*	34.39%*	22.08%*	42.18%*
AdaMo-CP [TA] [c1m]	40.49%	45.46%	28.36%	54.27%	26.61%	33.86%	21.84%*	41.58%*
AdaMo-CP [TA] [both]	40.51%	45.09%	28.07%	53.68%	27.03%	34.30%	22.25%	42.49%

we offer two different time budgets, the falling scores for the same environmental setting are marked with an asterisk.

Based on the experimental results shown in Table IV, the continual pretraining scheme shows poor performance on JAVA but could reliably improve the results on PYTHON.

If we compare the results of AdaMo-CP[DA] with AdaMo-CP[TA], then task-adaptive models usually perform more satisfying than domain-adaptive ones. Even though the domain-adaptive scheme owns the advantage in resource intensiveness, the task-adaptive one shows higher efficiency in data utilization. Therefore, we conclude that task-adaptive is more suitable for continual pretraining.

When other experimental settings are the same, pretraining only the decoder seems the optimal choice since the results of AdaMo-CP[DA][c1m] is better than those of other AdaMo-CP[DA] models, so does AdaMo-CP[TA][mlm]. A special case is AdaMo-CP[TA][mlm], which performs better on PYTHON.

The phenomenon can be explained by two reasons. First, code data follow strict grammar rules than text data; therefore the text data contain more extensive entropy, which indicates the decoder learning the interrelations of text tokens is more efficient. Second, when it is task-adaptive, the representation of PYTHON code could be easier to be optimized by the encoder.

As shown in Table V, if we assign a richer time budget for the continuous pretraining scheme, the patterns found before keep constant. Besides, all best results are contributed by CP models. Meanwhile, we found that a long training time seems not necessary, since almost half of the scores are reduced a bit. Overall, that continuously pretraining both encoder and decoder seldom promises better results than training only encoder or decoder. It should be caused by the incoordination of parallel representations and would be solved after sufficient finetuning.

RQ3 – Takeaway: The effects of continuous pretraining exist but vary with code data. Its combination with the task-adaptive scheme promises better improvements. It is recommended to only continuously pretrain the decoder since pretraining both the encoder and decoder is not necessarily better.

TABLE VI
RESULTS OF THE INTERMEDIATE FINETUNING SCHEME (24 HOURS)

Model	JAVA				PYTHON			
	C-BLEU	S-BLEU	METEOR	ROUGE	C-BLEU	S-BLEU	METEOR	ROUGE
AdaMo-basic	40.49%	45.30%	28.19%	53.99%	26.52%	33.85%	21.68%	41.25%
AdaMo-IF [DA] [CA]	29.06%	35.75%	20.53%	44.20%	25.60%	33.05%	21.02%	40.11%
AdaMo-IF [DA] [CE]	41.21%	46.33%	29.11%	55.59%	28.37%	35.44%	23.25%	44.28%
AdaMo-IF [DA] [CI]	40.65%	45.99%	28.83%	55.31%	27.66%	34.88%	22.96%	43.82%
AdaMo-IF [TA] [CA]	39.79%	44.63%	27.73%	53.19%	25.84%	33.42%	21.42%	40.97%
AdaMo-IF [TA] [CE]	40.50%	45.42%	28.02%	53.45%	26.53%	33.91%	21.15%	40.51%
AdaMo-IF [TA] [CI]	40.15%	45.33%	28.30%	54.24%	26.22%	33.65%	21.79%	41.44%

TABLE VII
SHOWCASE OF THE EFFECTS OF THE INTERMEDIATE TASKS

Sample	Model	Comment
Java#229	AdaMo-basic	returns true if the given word contains a whitespace
	AdaMo-IF [DA] [CE]	returns true if the input string contains a word except for word engines
	AdaMo-IF [DA] [CI]	check if the input string contains a word
	Ground Truth	returns true if the input string contains a word - breaking character
Java#547	AdaMo-basic	ensures that the object value is at the given allocation
	AdaMo-IF [DA] [CE]	ensures that a value is not null
	AdaMo-IF [DA] [CI]	ensures that the given location is not null
	Ground Truth	ensure the given value is not null and return it
Java#821	AdaMo-basic	draws a face of the specified shape
	AdaMo-IF [DA] [CE]	draws a circle for the given parameters
	AdaMo-IF [DA] [CI]	draws a dial on the given parameters
	Ground Truth	draws a cylinder for the given parameters
Python#201	AdaMo-basic	remove the most recent history from the history
	AdaMo-IF [DA] [CE]	remove all completed jobs
	AdaMo-IF [DA] [CI]	remove all the completed jobs from the history
	Ground Truth	remove all completed jobs from history
Python#439	AdaMo-basic	return a string description of the appropriate description on the given path
	AdaMo-IF [DA] [CE]	return a string describing the probable encoding of a file
	AdaMo-IF [DA] [CI]	return a string describing the probable encoding of a unicode path
	Ground Truth	return a string describing the probable encoding of a file
Python#660	AdaMo-basic	returns a set of all cliques of a chordal graph
	AdaMo-IF [DA] [CE]	return the set of maximal cliques of the chordal graph
	AdaMo-IF [DA] [CI]	returns set of maximal cliques of the chordal graph g
	Ground Truth	returns the set of maximal cliques of a chordal graph

D. Results of RQ4

The results for RQ4 are the summarization scores of ADAMO when applying the intermediate finetuning scheme, in comparison with our basic transfer learning approach.

According to the experimental results shown in Table VI, the combination of domain-adaptive finetuning with concept extrapolation achieves the best scores on all metrics. In contrast, the combination of domain-adaptive with concept annotation performs the worst and even worse than AdaMo-basic. All the other combinations have similar performance to AdaMo-basic, with slight improvements or a bit deteriorations.

When comparing domain-adaptive with task-adaptive, the former performs better than the latter if the intermediate stage task is CE or CI, but the situation becomes opposite if the stage task is CA. We believe that CE and CI are complex tasks while CA is rather a simple task. The domain-adaptive scheme learns from more data compared to the task-adaptive one, therefore it shows better generalization when given the same time budget. However, simple intermediate tasks might mislead the domain-adaptive scheme. The opposite case is that the performance of the task-adaptive scheme almost keeps unchanged no matter which stage task it is. Therefore, the domain-adaptive scheme should pair with complex intermediate tasks, while task-adaptive might not benefit too much from the intermediate finetuning.

Among all intermediate tasks, concept annotation is weaker than others. Meanwhile, CA is sensitive to the data in some cases, such as pairing with the domain-adaptive scheme on the JAVA data. However, other tasks usually promise better results and are insensitive to languages. On both JAVA and

TABLE VIII
OVERVIEW OF THE PROMOTION EFFECTS ON CODE SUMMARIZATION

Model	JAVA				PYTHON			
	C-BLEU	S-BLEU	METEOR	ROUGE	C-BLEU	S-BLEU	METEOR	ROUGE
AdaMo-basic	6.52%	0.76%	4.06%	0.84%	0.23%	0.12%	1.45%	0.17%
AdaMo-noise	6.87%	0.87%	4.28%	0.97%	1.28%	0.71%	2.71%	1.58%
AdaMo-CP [DA]	6.87%	0.58%	4.10%	0.75%	0.04%	0.06%	1.68%	0.56%
AdaMo-CP [TA]	6.84%	0.69%	4.06%	0.78%	2.42%	1.83%	4.12%	3.28%
AdaMo-IF [DA]	8.42%	3.05%	7.46%	3.83%	7.22%	4.82%	8.80%	7.53%
AdaMo-IF [TA]	6.55%	1.02%	4.47%	1.33%	0.26%	0.30%	1.97%	0.63%

PYTHON, CE and CI obtain at least similar but usually better results than AdaMo-basic, therefore, complex intermediate tasks are more stable and beneficial than simple ones. When pairing with the domain-adaptive scheme, concept extrapolation always performs better than concept interpolation, and while pairing the task-adaptive scheme, their results are close.

To check the effects of intermediate tasks on the generated comments, especially CE and CI, we could observe selected examples as shown in Table VII. The longest common substrings occurred in others that include the ground truth, but not in AdaMo-basic, are emphasized in italics, and they contribute the most to the improvements. Besides, the distinctions between CE and CI cause differences in other tokens, which are usually never appeared in the corresponding code data for the former CE case but often already appeared there for the latter CI case.

RQ4 – Takeaway: The combination of complex intermediate tasks with the domain-adaptive scheme is the best choice. The optimal intermediate task is concept extrapolation. The task-adaptive scheme promises more stable results, but its improvements are not so obvious as for the domain-adaptive scheme.

E. Discussion

To build a clear comparison of the effect promotion given by each extension, we compute the growth extents of best scores in each research question relative to those of SIT, as shown in Table VIII. Summarizing the results, we notice that the intermediate finetuning scheme pairing with domain-adaptive always boosts the results. On JAVA, the second largest growth origins from its combination with the task-adaptive way. On PYTHON, the second-best results are contributed by the continual pretraining scheme pairing with the task-adaptive way. Then, the refiner with adequate Gaussian noise performs the optimal while others have close results with the basic ADAMO.

It seems hard to understand that, when it is domain-adaptive, intermediate finetuning performs better than continuous pretraining. It seems that CP learns the latent representation by optimizing the self-supervised objectives, but IF learns the interrelations of tokens via our stage tasks in a supervised way. Therefore, CP is slower than IF when given the same time budget. When it is task-adaptive, CP and IF show their advantages on separate datasets. We believe it is due to the fact that the knowledge learned from the target dataset is more easily beneficial to the target task. Thereby, the results are likely to be intervened by data characteristics and fortuity.

As reflected in our results, Gaussian noise and continuous pretraining have a small contribution to the achieved effectiveness. There are no alternatives found yet to the AWGN emitter. We have experimented with other common noises in signal processing but only got negative results. For continuous pretraining, our results show that training the encoder did not help a lot, which indicates the MLM objective could be replaced with other token-level ones [40] for potential improvements.

F. Threats to Validity

Internal Validity. The most crucial limitation to our results comes from the representativeness of our evaluation data, although we already use the most common dataset for evaluation. In our early-stage experiments, we found that results for both baseline models and our transfer learning model, evaluated on the CSN data, are rather low, even though ours are still better. Based on our analysis to the phenomenon, we concluded that it is caused by the quality differences of code comments. Therefore, it would be meaningful to evaluate the performance of models on the data of various quality levels systematically.

External Validity. The results are limited in the way that they can be generalized to relevant generative tasks. Our approach applies to generative tasks where the input data and output data are sequential data, such as program migration and code generation, respectively from code to code and from text to code. The limitation is whether or not there are already pretrained models available for generating the data. However, the task that involves generating code is harder because the generated programs might be partial, buggy, and even specious. Therefore, there still exists potential challenges on how to effectively generalize our approach.

VI. RELATED WORK

In the following, we discuss prior works on code summarization and summarize the research status in topics of text summarization and program representation.

Code Summarization. In recent years, there have been several neural-based approaches proposed for code summarization. CODE-NN [41] uses the Long Short-Term Memory (LSTM) network [42] combined with global attention [43] for both code retrieval and code summarization. HYBRID-DRL [32] applies reinforcement learning to incorporate the AST structure and sequential content of snippets by using an actor-critic network.

Besides the token sequence, the parsing results of the code data, e.g., AST and CFG, are used for code summarization as well. HYBRID-DEEPCOM [30] fuses the lexical and syntactical information of code tokens and serialized ASTs using the Gated Recurrent Unit (GRU) network [44]. ATTENDGRU [45] applies the GRU encoders for code sequences and the serialized ASTs, and adopts attention components for their interrelations with the summary tokens.

Not only do neural models, but also information retrieval methods perform well in code summarization, therefore some work is inspired to combine them. RENCOS [46] utilizes the search engine to find the most semantically or syntactically

similar code snippets to augment samples. HYBRID GNN [47] constructs code property graph and meanwhile learns attention-based dynamic graph as the training data. In the process, it retrieves the most similar code and corresponding summary for information augmentation.

The latest models are focused on adapting Transformer, considering its success in natural language-related tasks. C2NL [48] is the enhanced Transformer [9] equipped with copy attention [49] and relative position encoding [50]. It merely relies on the knowledge from the sequence of code tokens. Furthermore, SIT [4] introduces the structure-induced attention mechanism to capture information from syntax structure, data flow, and data dependency. Follow the idea of building representations for AST [51], BASTS [3] decomposes the code data into blocks in CFG to generate a split AST for each of them and eventually generate corresponding representations. Therefore, it utilizes the syntax information from split ASTs, instead of the only original AST. Apart from these models, some work systematically studies current issues in code summarization, including the commonly used datasets and evaluation metrics [36], [37].

Text Summarization. Inspired by neural machine translation [21], [52], the sequence-to-sequence model with attention is proposed for abstractive summarization [53]. By leveraging the pointer network [54], PGNET [49] drives a pointer and a generator in parallel to freely choose from either the copied contents or the generated tokens. Meanwhile, it introduces the coverage mechanism [55] to penalize repetitions. Even further, BOTTOM-UP [56] first selects potential tokens for the summary and then generates the summary using the PGNET. SEQCOPYNET [57] extends the copy mechanism, which not merely learns to copy the isolated tokens, but also the subsequences. SAGCOPY [58] enhances the copy mechanism based on the token importance. It builds a directed graph and adopts the degree centrality to identify the key tokens.

Gradually, neural models equipped with copy mechanism are replaced by pretrained models, such as PEGASUS [59] for abstractive summarization, as well as MASS [60] and BART [61] for the general sequence-to-sequence tasks. Based on Transformer and transfer learning, universal models represented by T5 [62] are proposed, which are intended to solve most common NLP tasks at once. As the reflection of text summarization, SUMMEVAL [63] intends to resolve critical shortcomings in evaluation methods.

Program Representation. By simply regarding code data as token sequences, self-supervised representation learning could be tailored for code data, such as CODEBERT [11], CODEX [64] and PLBART [65]. With hypothesis that “*programs with the same functionality should have similar underlying representations*”, CONTRACODE [66] builds representations of program functionalities by learning from contrastive samples [67].

Also, it is common to first parse the code data into tree or graph structures for richer information. There have been many models proposed to learn the parsed results of code data. ASTNN [51] splits each AST into a sequence of small trees

for better representations. MRNCS [68] recaps serialization schemes on tree structures and categorized them into sampling-based [69] and traversal-based ones [31]. TDLS [70] utilizes GGNN [71] to learn both syntactic and semantic information. Opposite to static analysis, DYPRO [72] and LIGER [73] learn program representations through dynamic executions as well, from the mixture of symbolic and concrete execution traces.

VII. CONCLUSIONS

In this paper, we demonstrated the transfer learning model performs well in code summarization by assembling available foundation models, CODEBERT and GTP-2. Then we utilized Gaussian noise to optimize the latent representation of the assembly by simulating the context-aware settings. Last, we introduced continuous pretraining and intermediate finetuning as adaptive schemes for optional improvements. In addition, we proposed concept interpolation and concept extrapolation as the intermediate stage tasks for code summarization and validated their effectiveness. These tasks apply to general sequence-to-sequence learning as well. Moreover, we experimented with adaptation ideas to tailor foundation models with the data of related domains or designated tasks.

The goal of our work is to show that the transfer learning model based on existing foundation models is rather competitive and could even outperform SOTA models. Moreover, our results showed that neural models regarding code data merely as sequential data could still be powerful enough. Also, our model is compatible with various adaptive schemes, which promises further improvements in model performance. Compared with SOTA models, ADAMO is more friendly to potential rises in quality from either the model side or the data side.

Despite our results and findings, there are still questions waiting to be solved. For example, it is challenging to align the latent representations of pretrained encoder and decoder models efficiently, or else, we could take advantage of foundation language models just like playing building blocks. It might also be valuable to implement the refiner with certain neural models for the manipulation of latent representations.

REFERENCES

- [1] T. H. M. Le, H. Chen, and M. A. Babar, "Deep Learning for Source Code Modeling and Generation. Models, Applications, and Challenges," *Acm Computing Surveys*, vol. 53, no. 3, pp. 1–38, Jul. 2020.
- [2] Y. Zhu and M. Pan, "Automatic Code Summarization: A Systematic Literature Review," *arXiv:1909.04352 [cs.SE]*, 2019. [Online]. Available: <https://arxiv.org/abs/1909.04352>
- [3] C. Lin, Z. Ouyang, J. Zhuang, J. Chen, H. Li, and R. Wu, "Improving Code Summarization with Block-Wise Abstract Syntax Tree Splitting," *arXiv:2103.07845 [cs.SE]*, 2021. [Online]. Available: <https://arxiv.org/abs/2103.07845>
- [4] H. Wu, H. Zhao, and M. Zhang, "Code Summarization with Structure-Induced Transformer," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP*, 2021, pp. 1078–1090.
- [5] X. Han, Z. Zhang, N. Ding, Y. Gu, X. Liu, Y. Huo, J. Qiu, L. Zhang, W. Han, M. Huang *et al.*, "Pre-Trained Models: Past, Present and Future," *arXiv:2106.07139 [cs.AI]*, Aug. 2021. [Online]. Available: <https://arxiv.org/abs/2106.07139>
- [6] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill *et al.*, "On the Opportunities and Risks of Foundation Models," *arXiv:2108.07258 [cs.LG]*, 2021. [Online]. Available: <https://arxiv.org/abs/2108.07258>
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "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-HLT)*, 2019.
- [8] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving Language Understanding by Generative Pre-Training," OpenAI, Tech. Rep., 2018. [Online]. Available: <https://openai.com/blog/language-unsupervised/>
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," in *Conference on Neural Information Processing Systems (NIPS)*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., 2017, pp. 5998–6008.
- [10] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A Robustly Optimized BERT Pretraining Approach," *arXiv:1907.11692 [cs.CL]*, 2019. [Online]. Available: <https://arxiv.org/abs/1907.11692>
- [11] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang *et al.*, "CodeBERT: A Pre-Trained Model for Programming and Natural Languages," *arXiv:2002.08155 [cs.CL]*, 2020. [Online]. Available: <https://arxiv.org/abs/2002.08155>
- [12] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language Models Are Unsupervised Multitask Learners," Tech. Rep., 2019. [Online]. Available: <https://openai.com/blog/better-language-models/>
- [13] S. J. Pan and Q. Yang, "A Survey on Transfer Learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [14] S. Niu, Y. Liu, J. Wang, and H. Song, "A Decade Survey of Transfer Learning (2010–2020)," *IEEE Transactions on Artificial Intelligence*, vol. 1, no. 2, pp. 151–166, Oct. 2020.
- [15] T. George Karimpanal and R. Bouffanais, "Self-Organizing Maps for Storage and Transfer of Knowledge in Reinforcement Learning," *Adaptive Behavior*, vol. 27, no. 2, pp. 111–126, Dec. 2018.
- [16] S. Gururangan, A. Marasovi{Vc}, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, and N. A. Smith, "Don't Stop Pretraining: Adapt Language Models to Domains and Tasks," *arXiv:2004.10964 [cs.CL]*, 2020. [Online]. Available: <https://arxiv.org/abs/2004.10964>
- [17] Y. Pruksachatkun, J. Phang, H. Liu, P. M. Htut, X. Zhang, R. Y. Pang, C. Vania, K. Kann, and S. R. Bowman, "Intermediate-Task Transfer Learning with Pretrained Language Models: When and Why Does It Work?" in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020, pp. 5231–5247.
- [18] S. Storks, Q. Gao, and J. Y. Chai, "Recent Advances in Natural Language Inference: A Survey of Benchmarks, Resources, and Approaches," *arXiv:1904.01172 [cs.CL]*, 2019. [Online]. Available: <https://arxiv.org/abs/1904.01172>
- [19] S. Min, D. Chen, L. Zettlemoyer, and H. Hajishirzi, "Knowledge Guided Text Retrieval and Reading for Open Domain Question Answering," *arXiv:1911.03868 [cs.CL]*, 2019. [Online]. Available: <https://arxiv.org/abs/1911.03868>
- [20] S. Ruder, "Recent Advances in Language Model Fine-Tuning," 2021. [Online]. Available: <https://ruder.io/recent-advances-lm-fine-tuning/>
- [21] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to Sequence Learning with Neural Networks," in *International Conference on Neural Information Processing Systems (NIPS)*, 2014, pp. 3104–3112.
- [22] S. Rothe, S. Narayan, and A. Severyn, "Leveraging Pre-Trained Checkpoints for Sequence Generation Tasks," *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 264–280, Dec. 2020.
- [23] B. Li, H. Liu, Z. Wang, Y. Jiang, T. Xiao, J. Zhu, T. Liu, and c. li, "Does Multi-Encoder Help? A Case Study on Context-Aware Neural Machine Translation," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020, pp. 3512–3518.
- [24] J. Mallinson, A. Severyn, E. Malmi, and G. Garrido, "Felix: Flexible Text Editing Through Tagging and Insertion," *arXiv:2003.10687 [cs.CL]*, 2020. [Online]. Available: <https://arxiv.org/abs/2003.10687>
- [25] A. Aghajanyan, A. Gupta, A. Shrivastava, X. Chen, L. Zettlemoyer, and S. Gupta, "Muppet: Massive Multi-Task Representations with Pre-Finetuning," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021, pp. 5799–5811.
- [26] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in *Advances in Neural Information Processing Systems (NIPS)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlch{Vc}-Buc, E. Fox, and R. Garnett, Eds., 2019, pp. 8024–8035.
- [27] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Transformers: State-of-the-Art Natural Language Processing," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Oct. 2020, pp. 38–45.
- [28] S. Sharma, L. El Asri, H. Schulz, and J. Zumer, "Relevance of Unsupervised Metrics in Task-Oriented Dialogue for Evaluating Natural Language Generation," *arXiv:1706.09799 [cs.CL]*, 2017. [Online]. Available: <https://arxiv.org/abs/1706.09799>
- [29] H. Husain, H.-H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, "CodeSearchNet Challenge: Evaluating the State of Semantic Code Search," *arXiv:1909.09436 [cs.LG]*, 2019. [Online]. Available: <https://arxiv.org/abs/1909.09436>
- [30] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep Code Comment Generation with Hybrid Lexical and Syntactical Information," *Empirical Software Engineering*, vol. 25, no. 3, pp. 2179–2217, Jun. 2019.
- [31] —, "Deep Code Comment Generation," in *IEEE International Conference on Program Comprehension (ICPC)*, 2018, pp. 200–210.
- [32] Y. Wan, Z. Zhao, M. Yang, G. Xu, H. Ying, J. Wu, and P. S. Yu, "Improving Automatic Source Code Summarization Via Deep Reinforcement Learning," in *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2018, pp. 397–407.
- [33] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: A Method for Automatic Evaluation of Machine Translation," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2001, pp. 311–318.
- [34] A. Lavie and A. Agarwal, "METEOR: An Automatic Metric for MT Evaluation with High Levels of Correlation with Human Judgments," in *Workshop on Statistical Machine Translation (StatMT)*, 2007, pp. 228–231.
- [35] C.-Y. Lin, "ROUGE: A Package for Automatic Evaluation of Summaries," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2004, pp. 74–81.
- [36] D. Gros, H. Sezhiyan, P. Devanbu, and Z. Yu, "Code to Comment 'Translation': Data, Metrics, Baseline & Evaluation," in *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2020, pp. 746–757.
- [37] D. Roy, S. Fakhoury, and V. Arnaoudova, "Reassessing Automatic Evaluation Metrics for Code Summarization Tasks," in *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*, 2021, pp. 1105–1116.
- [38] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. Clement, D. Drain, D. Jiang, D. Tang *et al.*, "CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation," *arXiv:2102.04664 [cs.SE]*, 2021. [Online]. Available: <https://arxiv.org/abs/2102.04664>
- [39] Y. Xian, B. Schiele, and Z. Akata, "Zero-Shot Learning: The Good, the Bad and the Ugly," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 3077–3086.

- [40] S. Aroca-Ouellette and F. Rudzicz, "On Losses for Modern Language Models," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [41] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing Source Code Using a Neural Attention Model," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016, pp. 2073–2083.
- [42] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [43] T. Luong, H. Pham, and C. D. Manning, "Effective Approaches to Attention-Based Neural Machine Translation," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2015, pp. 1412–1421.
- [44] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning Phrase Representations Using RNN Encoder–Decoder for Statistical Machine Translation," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1724–1734.
- [45] A. LeClair, S. Jiang, and C. McMillan, "A Neural Model for Generating Natural Language Summaries of Program Subroutines," in *IEEE/ACM International Conference on Software Engineering (ICSE)*, 2019, pp. 795–806.
- [46] J. Zhang, X. Wang, H. Zhang, H. Sun, and X. Liu, "Retrieval-Based Neural Source Code Summarization," in *IEEE/ACM International Conference on Software Engineering (ICSE)*, 2020, pp. 1385–1397.
- [47] S. Liu, Y. Chen, X. Xie, J. Siow, and Y. Liu, "Retrieval-Augmented Generation for Code Summarization Via Hybrid GNN," in *International Conference on Learning Representations (ICRL)*, 2021.
- [48] W. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang, "A Transformer-Based Approach for Source Code Summarization," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020, pp. 4998–5007.
- [49] A. See, P. J. Liu, and C. D. Manning, "Get to the Point: Summarization with Pointer-Generator Networks," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017, pp. 1073–1083.
- [50] P. Shaw, J. Uszkoreit, and A. Vaswani, "Self-Attention with Relative Position Representations," in *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018, pp. 464–468.
- [51] J. Zhang, X. Wang, H. Zhang, H. Sun, K. Wang, and X. Liu, "A Novel Neural Source Code Representation Based on Abstract Syntax Tree," in *IEEE/ACM International Conference on Software Engineering (ICSE)*, 2019, pp. 783–794.
- [52] D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," in *International Conference on Learning Representations (ICRL)*, 2015.
- [53] A. M. Rush, S. Chopra, and J. Weston, "A Neural Attention Model for Abstractive Sentence Summarization," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2015, pp. 379–389.
- [54] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer Networks," in *International Conference on Neural Information Processing Systems (NIPS)*, 2015, pp. 2692–2700.
- [55] Z. Tu, Z. Lu, Y. Liu, X. Liu, and H. Li, "Modeling Coverage for Neural Machine Translation," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016, pp. 76–85.
- [56] S. Gehrmann, Y. Deng, and A. Rush, "Bottom-up Abstractive Summarization," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018, pp. 4098–4109.
- [61] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: Denoising Sequence-to-Sequence Pre-Training for Natural Language Generation, Translation, and Comprehension," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020, pp. 5858–5875.
- [57] Q. Zhou, N. Yang, F. Wei, and M. Zhou, "Sequential Copying Networks," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2018, pp. 4987–4994.
- [58] S. Xu, H. Li, P. Yuan, Y. Wu, X. He, and B. Zhou, "Self-Attention Guided Copy Mechanism for Abstractive Summarization," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020, pp. 1355–1362.
- [59] J. Zhang, Y. Zhao, M. Saleh, and P. J. Liu, "PEGASUS: Pre-Training with Extracted Gap-Sentences for Abstractive Summarization," in *International Conference on Machine Learning (ICML)*, 2020, pp. 11 328–11 339.
- [60] K. Song, X. Tan, T. Qin, J. Lu, and T.-Y. Liu, "MASS: Masked Sequence to Sequence Pre-Training for Language Generation," in *International Conference on Machine Learning (ICML)*, 2019, pp. 5926–5936.
- [62] C. Raffel, N. M. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," *arXiv:1910.10683 [cs.LG]*, 2020. [Online]. Available: <https://arxiv.org/abs/1910.10683>
- [63] A. R. Fabbri, W. Kryściński, B. McCann, C. Xiong, R. Socher, and D. Radev, "SummEval: Re-Evaluating Summarization Evaluation," *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 391–409, 2021.
- [64] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. Ponde, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, "Evaluating Large Language Models Trained on Code," *arXiv:2107.03374 [cs.LG]*, 2021. [Online]. Available: <https://arxiv.org/abs/2107.03374>
- [65] W. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang, "Unified Pre-Training for Program Understanding and Generation," in *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2021, pp. 2655–2668.
- [66] P. Jain, A. Jain, T. Zhang, P. Abbeel, J. Gonzalez, and I. Stoica, "Contrastive Code Representation Learning," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 5954–5971.
- [67] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality Reduction by Learning an Invariant Mapping," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 2006, pp. 1735–1742.
- [68] J. Gu, Z. Chen, and M. Martin, "Multimodal Representation for Neural Code Search," in *IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2021, pp. 483–494.
- [69] U. Alon, M. Zilberstein, O. Levy, and E. Yahav, "A General Path-Based Representation for Predicting Program Properties," in *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Jun. 2018, pp. 404–419.
- [70] M. Allamanis, M. Brockschmidt, and M. Khademi, "Learning to Represent Programs with Graphs," *arXiv:1711.00740 [cs.LG]*, 2018. [Online]. Available: <https://arxiv.org/abs/1711.00740>
- [71] Y. Li, D. Tarlow, M. Brockschmidt, and R. S. Zemel, "Gated Graph Sequence Neural Networks," *arXiv:1511.05493 [cs.LG]*, 2016. [Online]. Available: <https://arxiv.org/abs/1511.05493>
- [72] K. Wang, "Learning Scalable and Precise Representation of Program Semantics," *arXiv:1905.05251 [cs.PL]*, 2019. [Online]. Available: <https://arxiv.org/abs/1905.05251>
- [73] K. Wang and Z. Su, "Learning Blended, Precise Semantic Program Embeddings," *arXiv:1907.02136 [cs.SE]*, 2019. [Online]. Available: <https://arxiv.org/abs/1907.02136>