

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

12-2023

On the usage of continual learning for out-of-distribution generalization in pre-trained language models of code

Martin WEYSSOW

Xin ZHOU

Kisub KIM

David LO

Singapore Management University, davidlo@smu.edu.sg

Houari A. SAHRAOUI

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#), and the [Software Engineering Commons](#)

Citation

WEYSSOW, Martin; ZHOU, Xin; KIM, Kisub; LO, David; and SAHRAOUI, Houari A.. On the usage of continual learning for out-of-distribution generalization in pre-trained language models of code. (2023). *ESEC/FSE '23: Proceedings of ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, San Francisco, December 3-9*. 1470-1482.

Available at: https://ink.library.smu.edu.sg/sis_research/8574

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

On the Usage of Continual Learning for Out-of-Distribution Generalization in Pre-trained Language Models of Code

Martin Weyssow
DIRO, University of Montreal
Montreal, Canada
martin.weyssow@umontreal.ca

Xin Zhou
Singapore Management University
Singapore
xinzhou.2020@phdcs.smu.edu.sg

Kisub Kim*
Singapore Management University
Singapore
kisubkim@smu.edu.sg

David Lo
Singapore Management University
Singapore
davidlo@smu.edu.sg

Houari Sahraoui
DIRO, University of Montreal
Montreal, Canada
sahraouh@iro.umontreal.ca

ABSTRACT

Pre-trained language models (PLMs) have become a prevalent technique in deep learning for code, utilizing a two-stage pre-training and fine-tuning procedure to acquire general knowledge about code and specialize in a variety of downstream tasks. However, the dynamic nature of software codebases poses a challenge to the effectiveness and robustness of PLMs. In particular, world-realistic scenarios potentially lead to significant differences between the distribution of the pre-training and test data, *i.e.*, distribution shift, resulting in a degradation of the PLM's performance on downstream tasks. In this paper, we stress the need for adapting PLMs of code to software data whose distribution changes over time, a crucial problem that has been overlooked in previous works. The motivation of this work is to consider the PLM in a non-stationary environment, where fine-tuning data evolves over time according to a software evolution scenario. Specifically, we design a scenario where the model needs to learn from a stream of programs containing new, unseen APIs over time. We study two widely used PLM architectures, *i.e.*, a GPT2 decoder and a RoBERTa encoder, on two downstream tasks, API call and API usage prediction. We demonstrate that the most commonly used fine-tuning technique from prior work is not robust enough to handle the dynamic nature of APIs, leading to the loss of previously acquired knowledge *i.e.*, catastrophic forgetting. To address these issues, we implement five continual learning approaches, including replay-based and regularization-based methods. Our findings demonstrate that utilizing these straightforward methods effectively mitigates catastrophic forgetting in PLMs across both downstream tasks while achieving comparable or superior performance.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ESEC/FSE '23, December 3–9, 2023, San Francisco, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0327-0/23/12...\$15.00
<https://doi.org/10.1145/3611643.3616244>

CCS CONCEPTS

• Software and its engineering → Software libraries and repositories; • Computing methodologies → Natural language processing.

KEYWORDS

deep learning for code, pre-trained language models, continual learning, out-of-distribution generalization

ACM Reference Format:

Martin Weyssow, Xin Zhou, Kisub Kim, David Lo, and Houari Sahraoui. 2023. On the Usage of Continual Learning for Out-of-Distribution Generalization in Pre-trained Language Models of Code. 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/10.1145/3611643.3616244>

1 INTRODUCTION

Prior research [11, 19, 70] on code representation learning leverages a ubiquitous two-stage procedure to effectively train and specialize pre-trained language models (PLMs) for code-related downstream tasks. The first stage, *i.e.*, the pre-training, involves optimizing the model using self-supervised learning on a large dataset to acquire general knowledge about code. This pre-training phase allows the model to adapt to downstream tasks in the second stage, *i.e.*, the fine-tuning. Previous studies [1, 19, 72] typically leverage classical transfer learning methods, which consist of "transferring" the pre-trained knowledge to the target task by fine-tuning the model on a task-specific loss function and data. This approach has been successful in the fields of natural language processing (NLP) [8, 16] and deep learning for code [11, 19].

In this perspective, previous works [13, 63] have primarily focused on stationary settings, neglecting the practical need for models to adapt to changing environments and data over time. Most prior research [1, 19, 22] has suggested using transfer learning to fine-tune the model in static environments rather than addressing the dynamic nature of real-world scenarios. In practice, programming languages, software libraries and APIs are prone to change and evolution [28, 46, 48], leading to shifts in the distribution of the underlying software data over time, it is also known as concept drift [40, 65]. By ignoring the actual evolution of software codebases, existing studies [11, 63] have focused on fine-tuning




Fig. 1: Continual fine-tuning of a pre-trained language model of code. After pre-training, the model needs to adapt to new out-of-distribution (OOD) program data over time.

and testing pre-trained models of code using stationary datasets. In practice, the software evolution potentially leads to a noticeable difference between training and test data, *i.e.*, distribution shift, that is often not present in these stationary datasets. This phenomenon also occurs when the model is put into production and has to deal with real-world data [4, 26]. We argue that creating datasets that reflect real-world software evolution scenarios and distribution shifts is crucial in order to properly evaluate the **out-of-distribution (OOD) generalization** capability of code models [53]. The OOD generalization measures a model’s ability to generalize to new, unseen data with a significantly different distribution from the training data. Therefore, evaluating how PLMs of code generalize to OOD software data in software evolution scenarios appears as a prime issue.

Existing works on OOD generalization designed the datasets based on various distribution shifts in source code data [25, 30]. However, they did not address the problem of continually adapting a pre-trained model of code to streams of OOD data. The prime goal of our study is to explore methods for a model to better adapt to software evolution scenarios. In this context, we ask: *how to effectively continually fine-tune a pre-trained model of code to adapt to new data while still considering the past data?* (see Fig. 1). Over the past years, **continual learning (CL)** [47, 65] has emerged to address this problem, which is relevant to a wide range of research areas, including computer vision [6, 34, 38, 55] and NLP [7, 9, 56]. Although transfer learning methods are not tailored for continual learning scenarios, they can still operate to fine-tune a model on streams of data. However, these methods lack robustness, leading to unwanted phenomena such as forgetting past information, known as catastrophic forgetting [20, 43]. There exist other strategies, such as retraining the model from scratch using new data, which are also impractical due to the tremendous computational intensity in the pre-training phase. Motivated by these issues of the existing models, we attempt to investigate more robust and scalable fine-tuning techniques. We hypothesize that continual learning techniques may provide significant benefits over classical transfer learning in this context.

In this paper, we delve into the behavior of PLMs of code in a continual fine-tuning scenario, as depicted in Fig 1. Our objective is twofold: (1) to assess the out-of-distribution generalization capability of PLMs of code and (2) to investigate effective continual fine-tuning strategies to fine-tune the models in the presence of a stream of OOD data. Specifically, we address these challenges in a scenario reflecting how typical software codebases may evolve in

practice. To this end, we create five OOD domain datasets, each introducing new, unseen APIs by the models during their pre-training phase. These OOD datasets intend to simulate a stream of data for continual fine-tuning, and each dataset entails a significant distribution shift with respect to the pre-training data. As such, our setting establishes an OOD generalization problem. We consider two widely used model architectures: a GPT2-like [49] decoder and a RoBERTa-like [37] encoder pre-trained on code. To eliminate any data leakage between the pre-training and fine-tuning data, we decided to pre-train our models from scratch. We do not study the popular existing PLMs like CodeBERT [19] or CodeT5 [62] because they may be prone to potential data leakage, *i.e.*, seeing the OOD data in pre-training, that we cannot precisely control. We evaluate the models on two downstream tasks: API call prediction and API usage prediction. In the first task, the model attempts to predict API calls resulting in a single code token, given code tokens appearing before the call site. On the other hand, the second task involves the generation of the whole API usage resulting in a sequence of code tokens with the same input format as the prior task. Together, these two tasks provide a comprehensive evaluation of the model’s performance in different code generation scenarios.

We start by investigating the impact of OOD data on the performance of the GPT2-like decoder on both downstream tasks in a zero-shot setting, *i.e.*, without fine-tuning the model on the new OOD data. We find that the model consistently fails to generalize to OOD data by highlighting significant gaps in performance compared to in-distribution data across six evaluation metrics (*e.g.*, up to 75% drop in BLEU score). This finding strongly suggests that pre-training itself is not sufficient and cannot solve OOD generalization in PLMs of code. We then evaluate the models’ performance in the continual fine-tuning scenario using classical transfer learning and observe notable catastrophic forgetting. To address this issue, we implement a straightforward yet computationally inefficient cumulative fine-tuning approach by utilizing a replay buffer of infinite size. The results show that the approach drastically mitigates forgetting. Finally, we compare the performance of classical transfer learning to that of replay-based and regularization-based continual learning methods. Replay methods are considered tough-to-beat strategies for continual learning and consist of maintaining a small replay buffer containing samples from previously seen data. During fine-tuning, we use the replay buffer in conjunction with the current OOD training set to fine-tune the PLM. We explore regularization-based methods, including EWC [34], SI [71] and RWalk [10], which add regularization terms to the loss function at fine-tuning to prevent extensive changes in important parameters of the PLM. We chose those methods as they are computationally efficient, well-known, and considered strong baselines in the continual learning literature. We discover that those continual learning methods significantly reduce forgetting while achieving similar or superior effectiveness on both tasks.

To the best of our knowledge, this work constitutes the first initiative to study continual fine-tuning for OOD generalization of PLMs of code. We believe that the impact of continual learning in this research area has the potential to be far-reaching, particularly due to the inherent evolution of software data over time, and we discuss this aspect in more detail in the discussion section of the

paper (see Section 5). Our contributions can be summarized as follows:

- (1) We demonstrate that PLMs of code fail to generalize to OOD data and highlight the need for further investigation in this area.
- (2) We conduct a study on the behavior of two pre-trained model architectures of code in a continuous learning environment, showing that classical transfer learning lacks robustness and is prone to catastrophic forgetting.
- (3) We compare five continual learning methods, including replay-based and regularization-based approaches, in our continual fine-tuning scenario. We show the superiority of continual learning over classical transfer learning.
- (4) We provide a large-scale dataset of Java code snippets and their API usage sequences, including pre-training data and a procedure for extracting OOD data.

Organization. In Section 2, we discuss preliminaries on continual learning. In Section 3, we go through our experimental design. We present the results of our experiments in Section 4. In Section 5, we discuss the threats to the validity of our study, as well as potential broader impact and future research directions. We introduce the related work on out-of-distribution generalization and continual learning for pre-trained language models in Section 6. Finally, we discuss some future work and conclude this work in Section 7.

2 PRELIMINARIES ON CONTINUAL LEARNING

Existing PLMs such as BERT [16] or GPT [8] typically operate in transfer learning settings. By using a two-stage pre-training/fine-tuning procedure, these models can be specialized for a wide range of downstream tasks. However, in this setting, the data used for pre-training or fine-tuning are often assumed to be stationary, which is not reflective of real-world situations. In practice, transfer learning methods can still be applied to non-stationary data, such as a stream of data, but this technique is prone to catastrophic forgetting [20, 43].

To address the above issues, prior works [2, 20, 24, 34, 36, 59] introduced the concept of *continual learning* and designed specific techniques to mitigate catastrophic forgetting. The primary assumption for continual learning is that the neural network should possess the ability to adapt to new data or tasks while maintaining stability on previous data or tasks, often referred to as the plasticity–stability dilemma. Considering continual learning is particularly interesting for OOD generalization problems, as continual learning methods focus on a keeping good plasticity–stability trade-off. Altogether, it has to potential to enhance the generalizability of PLMs to a broader range of data. Continual learning methods often operate in constrained scenarios, and Hadsell et al. [24] outline a comprehensive list of objectives to balance in continual learning scenarios. There exist three main categories of methods for continual learning as defined in a previous study [15]. *Replay-based methods* store samples from previous experiences, *i.e.*, previous stream of data, in a replay buffer or use generative approaches to generate examples similar to those of previous experiences. The replay buffer is used in conjunction with the current experience data to train the model. Replay-based methods help the network gain stability by




Fig. 2: Procedure to extract the ID data used for model pre-training, and the OOD data used for continual fine-tuning.

enabling the network to train on previous samples, *i.e.*, stored in the replay buffer while adapting to new data. *Regularization-based methods* add a regularization term to the loss function to prevent catastrophic forgetting by penalizing changes to important neural network parameters. Examples of regularization-based methods include EWC [34], SI [71] and RWalk [10]. Finally, *parameter isolation methods* use dynamic architectures to incorporate knowledge from previous experiences to mitigate interference [52].

3 EXPERIMENTAL DESIGN

In this section, we describe the experimental setup of our study. We carefully control our data and model setup to implement our out-of-distribution scenario. We first outline the construction of our dataset and the generation of OOD data for continual fine-tuning. Next, we discuss the pre-training procedure of our models, the target downstream tasks and evaluation metrics. We present the results of our experiments in Section 4.

3.1 Dataset Construction

Pre-training language models from scratch require a large amount of data for the loss of the model to converge. With that in mind, we constructed our large dataset using programs crawled from GitHub using Google BigQuery¹. Specifically, we focused on Java programs and began by collecting all java files stored in GitHub repositories. Next, we used Groum [45] to extract all methods defined in the java files along with their API usage sequences. We extracted the API usage sequences to facilitate our data splitting and obtain the position of each API site inside the methods to implement our downstream tasks. Each sample consists of all the tokens of a method. To avoid duplication bias in our experiments [3], we deduplicated the dataset by comparing the hash of each method. The resulting dataset contains more than 68M Java methods. For our experiments, we shuffled these 68M methods and randomly selected 10M methods to constitute our initial dataset. Fig. 2 illustrates how we further split the data for our experiments. Because we chose the pre-train PLMs from scratch, we have to split our data into distribution (ID) data, used for model pre-training, and OOD data, used for continual fine-tuning. We also need to properly extract the OOD data to align with our scenario consisting of introducing new, unseen APIs over time to the PLM during fine-tuning.

¹<https://cloud.google.com/bigquery>

Table 1: Out-of-distribution dataset details.

DATASET	Domain	Package	Interfaces	# train	# test
\mathcal{D}_{OOD}^1	General	java.util.concurrent	BlockingQueue, ThreadPoolExecutor	47,213	5,239
		java.math	BigInteger		
		java.util	Base64, TreeSet		
		java.net	ForkJoinPool, Proxy, ServerSocket, SocketAddress, URLEncoder		
\mathcal{D}_{OOD}^2	Security	java.security	Cipher, CodeSource, Identity, KeyFactory, KeyPair, MessageDigest, Policy, Provider, Security, Timestamp	27,189	3,017
\mathcal{D}_{OOD}^3	Android	android.view	Display, InputEvent, Window	28,400	3,150
		android.widget	Checkbox, GridLayout		
		android.media	AudioFormat, ImageReader		
		android.hardware	Camera, Sensor		
		android.database	DatabaseUtils		
\mathcal{D}_{OOD}^4	Web	org.springframework	CacheManager, ClassPathResource, DataBuffer, HttpMessage, HttpRequest, JdbcTemplate, MessageChannel, MessageHandler, TaskExecutor	16,295	1,805
\mathcal{D}_{OOD}^5	Guava	com.google.common.graph	GraphBuilder, Network	13,448	1,489
		com.google.common.io	ByteSource, ByteStreams		
		com.google.common.cache	CacheBuilder, LoadingCache		
		com.google.common.collect	ListMultimap, Multimap		
		com.google.common.base	CharMatcher, Splitter		

Out-Of-Distribution Dataset – \mathcal{D}_{OOD} . We create five OOD datasets, $\mathcal{D}_{OOD}^1, \dots, \mathcal{D}_{OOD}^5$. Each OOD dataset represents a unique domain that encompasses a high-level functionality of APIs. For example, we have a domain *Security* that comprises APIs related to programming security-related code and a domain *Guava* that includes only APIs from the Guava² library. To create each OOD dataset, we randomly select 10 interfaces from packages/libraries related to their domain. Finally, we associate to each domain dataset all APIs within the selected interfaces, excluding class construction methods. Table 1 summarizes the dataset \mathcal{D}_{OOD} , which contains 147,245 samples in total.

To form each OOD dataset, we select samples from the pool of 10 million Java methods that manipulate at least one of their associated API. In our experiments, we perform continual fine-tuning on the training sets associated with the OOD dataset $\mathcal{D}_{OOD}^1, \dots, \mathcal{D}_{OOD}^5$ sequentially. Therefore, to prevent data leakage, we exclude samples that manipulate APIs from multiple domains. This elimination of samples removes a significant threat to the validity of our OOD scenario and ensures that APIs are introduced as intended during the fine-tuning process. To obtain representative test sets, we randomly select 10% of samples that manipulate each API within each OOD dataset and used the selected samples to form the corresponding domain test set.

In-Distribution Dataset – \mathcal{D}_{ID} . We obtain \mathcal{D}_{ID} by removing the samples in \mathcal{D}_{OOD} from the initial data. Then, we shuffle \mathcal{D}_{ID} and randomly select 50,000 samples for test (\mathcal{D}_{ID_test}). \mathcal{D}_{ID_PT} contains the remaining samples for pre-training, and we randomly select 100,000 for model validation ($\mathcal{D}_{ID_PT_valid}$). In particular, those samples allow us to monitor the evolution of the loss of the model on an independent validation set to avoid overfitting the pre-training data. In total, the pre-training set $\mathcal{D}_{ID_PT_train}$ contains more than 9M samples to pre-train the models.

²<https://github.com/google/guava>

3.2 Models and Tasks Setup

In this work, we consider two widely-used deep learning architectures for code: a RoBERTa-like encoder [37] and a GPT2-like decoder [49]. We deliberately exclude the utilization of large language models (LLMs) in our research due to the substantial computational resources essential for their pre-training. To comprehensively address our OOD scenario, it is imperative to pre-train a model from scratch prior to continually fine-tune it on code containing new, unseen APIs. Consequently, we opt to evaluate two smaller models architectures, namely RoBERTa and GPT-2, which either serve as foundational models for PLMs like CodeBERT [19] or to generative models.

Decoder – \mathcal{M}_{dec} . The decoder model is based on the GPT-2 architecture, with the same hyperparameters, and is pre-trained using a causal language modeling objective, *i.e.*, left-to-right next token prediction. As we conducted our experiments under limited resources, we implemented a small version of GPT-2 with 110 million trainable parameters and pre-train the model for 100,000 steps. We use early stopping to select the best model checkpoint, based on the loss on the validation set $\mathcal{D}_{ID_PT_valid}$.

Encoder – \mathcal{M}_{enc} . The encoder model is based on the RoBERTa architecture, with the same hyperparameters, and is pre-trained using a masked language modeling objective. We implemented a base version of RoBERTa. The model has 125 million trainable parameters and is pre-trained similarly to the decoder model, with early stopping used to select the best checkpoint. Note that conversely to \mathcal{M}_{dec} , the encoder’s architecture is not suitable for generation tasks. Therefore, we add a randomly initialized language modeling head on top of it for fine-tuning using the OOD datasets. As a result, we expect \mathcal{M}_{enc} to be less stable than \mathcal{M}_{dec} and more prone to catastrophic forgetting since the language modeling head is not pre-trained. This comparison provides valuable insights into the robustness of two different architectures.

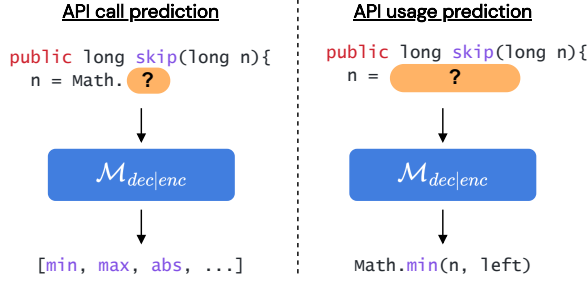


Fig. 3: Overview of the downstream tasks. In the API call prediction task, the model outputs a list of top- k candidates to predict the API call token (i.e., min). In the API usage prediction task, the model attempts to predict all the tokens constituting the API usage (interface name, method name, parameters and syntactical tokens). The models only leverage left-context tokens to generate a prediction.

Downstream Tasks. We employ two downstream tasks to evaluate the ability of our PLMs of code to learn and adapt to new software data that introduce new, unseen APIs over time. Fig. 3 illustrates both tasks. For API call prediction, the model takes as input all the tokens of the method preceding the call site of the API and generates top- k candidates. For API usage prediction, the model takes as input the same tokens as for the API call prediction task, but attempts to generate the whole API usage (interface name, method name, parameters and syntactical tokens), which constitutes a more challenging task. The rationale for evaluating the PLMs on these two downstream tasks is to select tasks where prior knowledge about the APIs seems decisive to effectively perform the task. Consequently, the choice for these two tasks is highly relevant to our continual OOD scenario, and it allows us to directly measure the impact of OOD APIs on the effectiveness of the PLMs. In Section 5.2, we discuss the applicability of our methodology to other code-related tasks.

Evaluation Metrics. We measure the performance of the models on both downstream tasks with metrics used in prior works. For API call prediction, we report the Exact Match@ k (EM@ k), which gives the percentage of correct predictions when considering lists of k candidates. For API usage prediction, we report BLEU score, Exact Match (EM), and CodeBLEU [50].

To measure how the models perform in a continual learning environment, we use two meta-metrics adapted from prior works [10, 32]: the *Average (A)* and *Forgetting (F)* metrics. We define the average A_M of a metric M on a test dataset \mathcal{D}_{OOD}^i as:

$$A_M = \frac{1}{T} \sum_{j=i}^T M_j(\mathcal{D}_{OOD}^i),$$

where j refers to the next incremental learning steps after the i -th included. M_j denotes an evaluation metric, e.g., EM@ k , computed at time step j on the test set and T denotes the maximum number of fine-tuning steps, i.e., five in our case. The Average metric only gives information on how accurate the model is but does not provide any insight into its ability to mitigate catastrophic forgetting. We define

Table 2: API call prediction results in zero-shot using \mathcal{M}_{dec} .

DATASET	METRICS		
	EM@1	EM@5	EM@10
\mathcal{D}_{ID_test}	<u>72.88</u>	<u>83.30</u>	<u>85.60</u>
\mathcal{D}_{OOD}	40.82 (44% ↓)	51.19 (38.5% ↓)	54.17 (36.7% ↓)
\mathcal{D}_{OOD}^1	49.91 (31.6% ↓)	62.0 (25.6% ↓)	64.46 (24.6% ↓)
\mathcal{D}_{OOD}^2	53.72 (26.3% ↓)	62.59 (24.8% ↓)	64.93 (24.2% ↓)
\mathcal{D}_{OOD}^3	23.78 (67.4% ↓)	32.64 (60.8% ↓)	36.33 (57.6% ↓)
\mathcal{D}_{OOD}^4	30.72 (57.9% ↓)	43.67 (47.3% ↓)	47.89 (44% ↓)
\mathcal{D}_{OOD}^5	37.54 (48.6% ↓)	49.53 (40.6% ↓)	53.22 (47.9% ↓)

the forgetting F_M^k of a metric M on a test dataset \mathcal{D}_{OOD}^i at time step k as:

$$F_M^k = M_i(\mathcal{D}_{OOD}^i) - M_k(\mathcal{D}_{OOD}^i), \quad i < k.$$

This is the difference between the first time the metric is computed, i.e., after fine-tuning the model on \mathcal{D}_{OOD}^i at time step i , and the metric computed at time step k . F_M^k gives information on the stability of the model, i.e., its capability to not forget from the past. Therefore, the lower F_M^k , the better.

Implementation Details. To pre-train \mathcal{M}_{dec} and \mathcal{M}_{enc} , we used four Tesla V100-SXM2-32GB GPUs. It took about 7 days to pre-train \mathcal{M}_{dec} , and 2 days to pre-train \mathcal{M}_{enc} . For fine-tuning and inference, we used a single Tesla V100-SXM2-32GB GPU. We used Huggingface’s libraries [67] to implement the models and store the datasets. To implement the continual learning approaches, we used Avalanche [39]. We provide all the implementation details of our experiments and release our data publicly in our replication package (see Data Availability section).

4 EXPERIMENTAL RESULTS

4.1 How Does \mathcal{M}_{dec} Generalize to ID and OOD Data in Zero-Shot?

In this experiment, we evaluate the performance of the model \mathcal{M}_{dec} on the ID and OOD test data in a zero-shot setting for both downstream tasks. We do not experiment with \mathcal{M}_{enc} as the model is not capable of generating code before fine-tuning and, therefore, cannot operate in a zero-shot setting. The purpose of this experiment is twofold. First, it aims to validate the experimental setup of our study. If we observe significant differences in the evaluation metrics obtained on the ID and OOD datasets, it would suggest that our OOD scenario is well-formed and reasonable. Secondly, significant gaps between the ID and OOD test data imply that PLMs such as \mathcal{M}_{dec} still require the use of robust transfer learning or continual learning techniques to generalize to new data without forgetting about past data.

API Call Prediction. Table 2 reports the EM@1, EM@5 and EM@10 on the ID and OOD test datasets. The results show that the model performs well on ID data, reaching almost 73% in EM@1. However, when tested on OOD data, the performance drops significantly. The decline in performance is less severe when considering more

Table 3: API usage prediction results in zero-shot using \mathcal{M}_{dec} .

DATASET	METRICS		
	BLEU	EM	CodeBLEU
\mathcal{D}_{ID_test}	<u>21.19</u>	<u>51.54</u>	<u>29.94</u>
\mathcal{D}_{OOD}	8.57 (59.56% ↓)	33.74 (34.54% ↓)	20.03 (33.10% ↓)
\mathcal{D}_{OOD}^1	5.94 (71.97% ↓)	34.29 (33.47% ↓)	15.71 (47.53% ↓)
\mathcal{D}_{OOD}^2	11.81 (44.27% ↓)	40.46 (21.50% ↓)	25.64 (14.36% ↓)
\mathcal{D}_{OOD}^3	7.26 (65.74% ↓)	28.01 (45.65% ↓)	16.49 (44.92% ↓)
\mathcal{D}_{OOD}^4	15.55 (26.62% ↓)	29.39 (42.98% ↓)	19.72 (34.13% ↓)
\mathcal{D}_{OOD}^5	5.11 (75.88% ↓)	30.71 (40.42% ↓)	25.81 (13.79% ↓)

API call candidates, but it remains a significant issue. Furthermore, variations in the performance decline are observed across different OOD datasets. For example, the model performs better on the Security domain (\mathcal{D}_{OOD}^2) than domains such as Android (\mathcal{D}_{OOD}^3) or Web (\mathcal{D}_{OOD}^4), which likely contain more domain-specific API calls.

API Usage Prediction. Table 3 reports the BLEU score, EM and CodeBLEU score on both ID and OOD test datasets. The results indicate that the model performs poorly on OOD data in comparison to ID data, with significant decreases in all evaluation metrics. Additionally, we notice that the EM and CodeBLEU metrics vary similarly to the EM@k metrics on the API call prediction task. The Android and Web domains experience the most severe drops, whereas the Security domain experiences the least severe drop.

Our results demonstrate that the model \mathcal{M}_{dec} (without fine-tuning) is unable to generalize to OOD data while showing strong performance on ID data. Our findings also support the validity of our OOD dataset as a realistic and meaningful test of the model’s ability to adapt to new data in a continuous environment.

4.2 Do Models Forget About Past Data Using Classical Transfer Learning?

In this section, we evaluate how classical transfer learning, *i.e.*, using fine-tuning as in prior work, performs in the continual learning scenario. We fine-tune the models \mathcal{M}_{dec} and \mathcal{M}_{enc} sequentially on the stream of OOD datasets $\mathcal{D}_{OOD}^1, \dots, \mathcal{D}_{OOD}^5$. We refer to this approach as “naive fine-tuning”, a common term used in the continual learning literature to refer to classical transfer learning, as it does not utilize mechanisms to address catastrophic forgetting. We report the results in terms of EM@1 for API call prediction and EM for API usage prediction. Fig. 4 illustrates the evolution of the EM@1 and EM metrics on the OOD test sets throughout the fine-tuning steps for both models. Each column of a heatmap refers to the evolution of the performance of the model on a particular test set, and each row refers to a new incremental fine-tuning step. Note that we do not compute the metric on a test set whose corresponding training set has not been seen yet by the model. To quantify catastrophic forgetting, we report the Forgetting (F) metrics of the EM@1 and EM metrics in Table 4. We do not report all the values

Table 4: Forgetting metrics for the naive fine-tuning baseline.

MODEL	DATASET	$F_{EM@1}^5$	F_{EM}^5
\mathcal{M}_{dec}	General ($\Delta t = 4$)	5.64	13.00
	Security ($\Delta t = 3$)	6.71	13.55
	Android ($\Delta t = 2$)	6.77	10.68
	Web ($\Delta t = 1$)	1.80	5.09
\mathcal{M}_{enc}	General ($\Delta t = 4$)	10.99	11.80
	Security ($\Delta t = 3$)	23.38	22.74
	Android ($\Delta t = 2$)	11.15	11.91
	Web ($\Delta t = 1$)	10.99	7.23

for every previously introduced metric as we have a strict page limit, and report them in our replication package.

Fine-Tuning Details. At each time step t , we fine-tune the models’ checkpoints from the previous time step on the dataset \mathcal{D}_{OOD}^t . We select 10% of the training samples from each OOD dataset as a validation set. For each fine-tuning, we set the number of training epochs to 10 and use early stopping by monitoring the evolution of the validation loss with a patience of two epochs. We keep the best checkpoints of the models at each fine-tuning step t and compute the task metrics on the previous and current test sets.

Performance of \mathcal{M}_{dec} and \mathcal{M}_{enc} . In Fig. 4, each heatmap depicts the evolution of a metric on the test sets for a single model on one task. The diagonal values in the heatmaps indicate the metric computed on the test set of the current OOD dataset. We observe substantial catastrophic forgetting for both tasks and models and all domains and metrics. That is, we observe a decline of the metrics in all columns, indicating that the model forgets the previous domains when fine-tuned on a new domain. For example, the EM@1 on \mathcal{D}_{OOD}^1 (General) drops from 57.37% to 51.73% for \mathcal{M}_{dec} . Another example, is the EM on \mathcal{D}_{OOD}^2 (Security) dropping from 40.79% to 18.05% for the model \mathcal{M}_{enc} . A glance at the heatmaps suggests that the forgetting is more severe for the encoder \mathcal{M}_{enc} . Overall, as we increase the number of fine-tuning steps, the forgetting further intensifies in most cases. In addition, for the decoder, the decline in the metrics after one fine-tuning step is less significant compared to the encoder. For example, after one fine-tuning step, the EM@1 on \mathcal{D}_{OOD}^2 drops from 60.93% to 57.66% (−3.27%) for the decoder. Whereas it drops from 58.37% to 32.94% (−25.43%) for the encoder. This means that more fine-tuning steps are required for the decoder to forget about past data more severely, whereas, for the encoder, one fine-tuning step is already enough to show a significant decline in performance. This observation confirms our intuition expressed in Section 3.2 that \mathcal{M}_{enc} may be less stable than \mathcal{M}_{dec} due to the additional language modeling head randomly initialized.

Forgetting Metrics. In Table 4, we calculate the Forgetting metric for the EM@1 and EM metrics and for both models. Note that we calculate the F metric at the final time step of the continual fine-tuning. According to the heatmaps of Fig. 4, the F^5 metric of a domain is the difference between the first and last value of its corresponding column. This difference represents the amount of forgetting that has occurred on each OOD domain during fine-tuning. The Δt in the table indicates how recently the model was

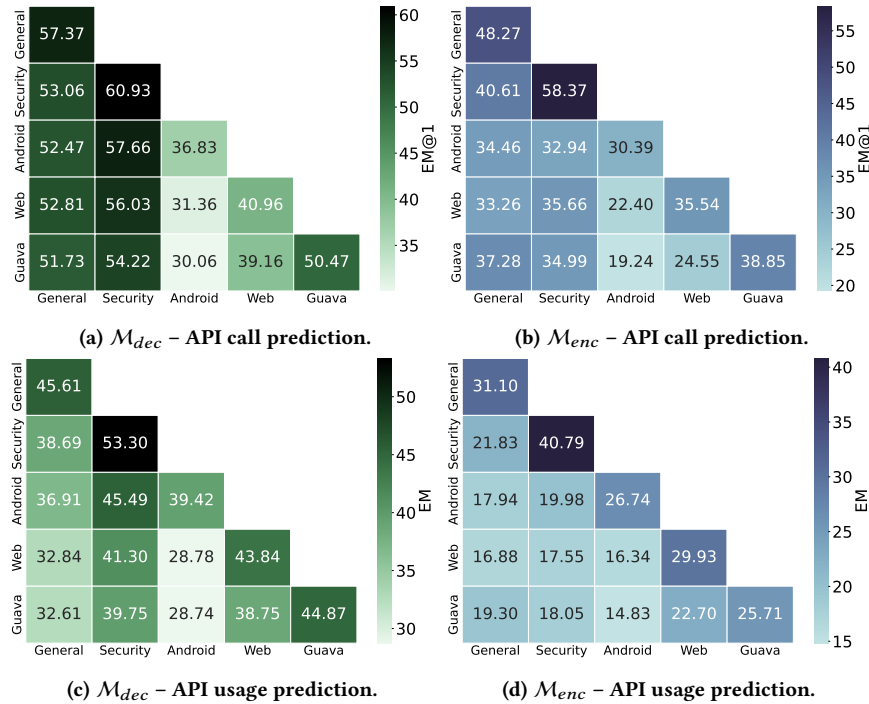


Fig. 4: Naive fine-tuning approach results.

fine-tuned on a particular domain dataset. We notice that for the decoder M_{dec} , the forgetting is less severe for the EM@1 (used in the API call prediction) than for the EM (used in the API usage prediction). The difference can be attributed to the fact that the API call prediction task is substantially easier than the API usage prediction task. In general, we observe more severe forgetting for the encoder, which further confirms our intuition about the lack of stability of M_{enc} .

Our results and observations illustrate that the problem of forgetting about past data is a major issue for both studied models and significantly more severe for the model M_{enc} . Even with a low number of fine-tuning steps, catastrophic forgetting is already prominent. By considering more fine-tuning steps, we can expect the problem to exacerbate.

We conclude that classical transfer learning, the most commonly used fine-tuning method in prior work, is not sufficient and robust enough to allow the model to adapt to new data while retaining knowledge of past data.

4.3 How Do Continual Learning Approaches Compare to Classical Transfer Learning?

To tackle the problem of catastrophic forgetting highlighted in our previous experiments, we propose to leverage some commonly used continual learning approaches from the literature. In this experiment, the naive fine-tuning approach is the lower-bound baseline, as it has no designed mechanism to mitigate catastrophic forgetting. We begin by introducing an upper-bound approach, referred to as




Fig. 5: Comparison of naive and cumulative fine-tuning settings for both models on API call prediction (EM@1).

"cumulative fine-tuning", which involves storing all training samples from each OOD training set cumulatively. With this approach, we perform continual fine-tuning using all samples from previous fine-tuning steps in addition to the current ones. This approach is usually upper-bound in continual learning settings as by storing all samples from previous data, the model can optimize its learning to generalize better to the whole stream of data. However, the cumulative fine-tuning approach is not usable in practice for a couple of reasons: (1) we may not always have access to all previous data at any given time, and (2) it requires storing all previous samples and significantly more computations during fine-tuning. This upper-bound approach aims to minimize forgetting while achieving the




Fig. 6: Comparison of naive and cumulative fine-tuning settings for both models on API usage prediction (EM).

best overall performance. We compare the cumulative and naive approaches in Fig. 5 and Fig. 6. Next, we introduce additional CL methods, including a replay-based method and three regularization-based methods: EWC [34], SI [71], and RWalk [10]. One advantage of these three methods over the replay method is that they do not require storing samples from previous data while fine-tuning. We report the Average (A) and Forgetting (F) metrics for both tasks and models on the EM@1 and EM metrics in Table 5 and Table 6. Note that there is no Forgetting metric for Guava as it is the last domain the PLMs are fine-tuned on.

Fine-Tuning Details. We use the same fine-tuning procedure as in the previous experiment. For the replay baseline, we set the buffer size to 200, *i.e.*, number of samples stored from past OOD training sets. We provide all our hyperparameters and further details about the implementations in our replication package.

Cumulative Fine-Tuning. In Fig. 5, we compare the naive and cumulative approaches for the API call prediction task ($EM@1$) on both decoder and encoder models. Each curve illustrates the evolution of the EM@1 on a particular OOD test set. The figure further demonstrates how the naive approach (bottom-left part of the figure) with the encoder leads to significantly more forgetting than for the decoder, as previously discussed. At the left of Fig. 5, we observe that the cumulative fine-tuning approach effectively eliminates the catastrophic forgetting issue for both models. Specifically, the EM@1 does not decrease over time and even increases throughout the fine-tuning, indicating improvement during continual fine-tuning, also known as positive transfer. In Fig. 6, we make the same observations for the API usage prediction task on the EM metric.

Continual Learning Approaches. Table 5 reports the Average and Forgetting metrics of the EM@1 on each OOD test set for \mathcal{M}_{dec} and \mathcal{M}_{enc} , with the naive fine-tuning approach as baseline. Similarly to Section 4.2, we compute the F metric at the end of the continual fine-tuning. Firstly, we observe that for both models, the cumulative fine-tuning approach is the best option to mitigate catastrophic forgetting and generally leads to the best $A_{EM@1}$. With the cumulative approach, the $F_{EM@1}^5$ metric is always negative,

which indicates a positive transfer (an increase in the EM@1). For instance, we get -8.02 in $F_{EM@1}^5$ for \mathcal{M}_{dec} in the Security domain, *i.e.*, an increase of $+8.02$ in the metric through fine-tuning. However, we observe large gaps between the $A_{EM@1}$ obtained using the cumulative approach and the naive approach on the Guava dataset (last fine-tuning step). We hypothesize that with an ever-increasing replay buffer, the models can no longer learn from new data and thus lose their ability to adapt with time. In addition to being computationally intensive, the cumulative fine-tuning approach is not scalable and robust, as previously mentioned. Overall, all other CL approaches, except EWC, greatly reduce forgetting and show a superior average EM@1 compared to the naive approach. The Replay approach generally produces the best or second best $A_{EM@1}$. Without the cumulative approach, RWalk is the best method to mitigate forgetting for \mathcal{M}_{dec} , whereas SI is better for \mathcal{M}_{enc} . In Table 6, we report the results for the API usage prediction task. We observe similar trends, except that the Replay approach is less effective for both models. However, RWalk and SI are the best methods for \mathcal{M}_{dec} and \mathcal{M}_{enc} , respectively.

In this final experiment, we demonstrate that continual learning methods, including two replay-based methods (Replay and Cumulative) and two regularization-based methods (SI and RWalk) effectively reduces catastrophic forgetting while achieving similar or superior effectiveness compared to classical transfer learning on both tasks.

5 DISCUSSION

In this section, we address some threats to the validity of our study. We then discuss the broader impact of our study and various opportunities for future work.

5.1 Threats to Validity

Threats to External Validity. We identified a main threat regarding the monolingual aspect of our dataset. Our OOD scenario requires extracting API usage sequences from the source code. Therefore, integrating more programming languages demands substantial additional effort, which we deliberately leave for future work. In addition, the construction of our dataset does not include any programming language-specific design and avoids any data leakage between the ID and OOD data. Consequently, it is highly likely that our results are not affected by the programming language of the data.

Another threat related to the data is the choice of the OOD domains and APIs. To mitigate this threat, we selected five domains covering different types of programs. Specifically, we selected 10 random interfaces per domain. Our results show that catastrophic forgetting is observed consistently for all domains, and the selection of different interfaces would result in different intensities in forgetting. We leave the study of this qualitative aspect for future work.

The choice of the downstream tasks presents another external threat to validity of our study. We employed two generation tasks, API call and API usage prediction. We focus on APIs-related tasks because APIs are an important part of the distribution of code tokens in programs and give lots of information about the semantics

Table 5: Continual learning approaches results for API call prediction using the EM@1 metric.

MODEL	METHOD	General		Security		Android		Web		Guava	
		$A_{EM@1} \uparrow$	$F_{EM@1}^5 \downarrow$	$A_{EM@1}$	$F_{EM@1}^5$	$A_{EM@1}$	$F_{EM@1}^5$	$A_{EM@1}$	$F_{EM@1}^5$	$A_{EM@1}$	$F_{EM@1}^5$
\mathcal{M}_{dec}	Naive	53.49	5.64	57.21	6.71	32.75	6.77	40.06	1.80	50.47	-
	EWC [34]	53.22	7.02	57.16	7.49	33.73	5.72	40.14	3.77	49.59	-
	SI [71]	54.65	3.57	59.24	3.45	34.04	2.39	38.93	1.36	48.16	-
	RWalk [10]	54.38	2.39	57.39	2.80	31.64	1.97	38.19	1.65	45.28	-
	Replay	55.66	4.41	58.87	2.98	34.66	2.01	41.12	2.41	49.72	-
	Cumulative	55.63	-0.51	58.44	-8.02	35.74	-0.73	32.99	-3.01	42.79	-
\mathcal{M}_{enc}	Naive	38.78	10.99	40.49	23.38	24.01	11.15	30.05	10.99	38.85	-
	EWC [34]	39.38	9.84	44.10	22.15	23.93	10.58	29.22	7.53	40.66	-
	SI [71]	44.29	5.94	50.05	8.10	21.39	4.02	27.79	2.56	35.67	-
	RWalk [10]	43.42	6.07	48.05	14.74	22.23	7.10	29.75	4.37	36.10	-
	Replay	45.15	5.48	51.56	10.56	24.31	8.27	32.53	3.92	40.22	-
	Cumulative	48.06	-0.92	56.40	-3.15	29.59	-3.62	27.79	-1.65	33.10	-

Table 6: Continual learning approaches results for API usage prediction using the EM metric.

MODEL	METHOD	General		Security		Android		Web		Guava	
		$A_{EM} \uparrow$	$F_{EM}^5 \downarrow$	A_{EM}	F_{EM}^5	A_{EM}	F_{EM}^5	A_{EM}	F_{EM}^5	A_{EM}	F_{EM}^5
\mathcal{M}_{dec}	Naive	37.32	13.00	44.96	13.55	32.31	10.68	41.30	5.09	44.87	-
	EWC [34]	36.88	12.95	44.84	13.08	33.92	9.46	39.00	6.73	45.71	-
	SI [71]	40.36	8.26	49.58	6.89	30.01	3.24	36.95	1.65	43.14	-
	RWalk [10]	40.43	6.23	47.11	4.04	33.34	2.63	36.54	2.13	41.22	-
	Replay	39.49	11.11	46.88	8.21	33.39	7.63	39.49	6.08	43.65	-
	Cumulative	43.29	2.02	47.26	-13.33	36.09	-2.28	27.92	-4.59	31.35	-
\mathcal{M}_{enc}	Naive	21.41	11.80	24.09	22.74	19.30	11.91	26.32	7.23	25.71	-
	EWC [34]	21.32	11.53	26.36	21.02	19.43	11.96	25.74	8.38	28.74	-
	SI [71]	27.22	5.03	30.85	8.23	18.57	2.20	23.03	1.65	21.26	-
	RWalk [10]	25.21	8.80	29.25	12.23	19.10	7.62	25.00	4.28	24.23	-
	Replay	23.48	13.54	29.94	13.96	18.09	11.88	24.51	5.92	26.48	-
	Cumulative	30.50	3.05	35.89	-6.88	24.81	-4.88	21.88	-1.97	18.43	-

of programs. We observe significant catastrophic forgetting in these two API-related tasks and hypothesize that catastrophic forgetting could appear in other SE tasks because of the importance of APIs in code. For instance, previous work found that APIs play important roles in writing the summarization of code [31], detecting code clones [44], retrieving code given a query [42], etc. We leave the investigation of the OOD phenomenon in other tasks as future work.

We identified an external threat to validity related to the limited number of fine-tuning steps in our continual fine-tuning settings. In practice, a PLM deployed to a real production environment would potentially face a larger number of fine-tuning steps throughout its lifetime. In this paper, we showed that both PLMs suffer from severe catastrophic forgetting, although we only consider five fine-tuning steps. We also demonstrated that more steps generally result in more forgetting about past data.

Finally, the selection of the size of the PLMs, in terms of the number of trainable parameters, constitutes a potential threat to the validity of our study. While increasing the number of parameters may still result in OOD generalization issues due to the design of our datasets, it is uncertain whether catastrophic forgetting would occur with the same magnitude for larger models. Our experiments were performed under limited computational resources, which required us to consider architectures with a limited number

of parameters. To mitigate this threat, we maximized the size of the models considering our limited resources. We pre-train PLMs with 110M and 125M parameters which are within the range of PLMs such as CodeBERT [19], CodeT5 [62] or CodeGPT [41].

Threats to Internal Validity. The hyperparameter choices for our CL approaches constitute the main threat to internal validity. We selected our hyperparameters based on values used in prior works about continual learning [10, 32, 34, 71]. These hyperparameters can be optimized for our scenario by using search methods, which tend to have a high computational cost. However, this aspect is not critical to the study as we have already shown the advantages of incorporating continual learning techniques with reasonable hyperparameter values.

Threats to Construct Validity. We identified one threat to construct validity related to the choice of our evaluation metrics. We mitigate this threat by selecting metrics widely used in prior works to evaluate code generation tasks [50, 68]. Additionally, we adapted continual learning metrics from prior works [10, 32] to evaluate our continual fine-tuning scenario.

5.2 Broader Impact and Opportunities

Our study sheds light on the performance of PLMs of code in a continual learning setting for out-of-distribution generalization.

We believe that this initial exploration of continual learning for code (*CL4Code*) will inspire further investigation in this important area. Our findings highlight two potential areas for future research: improving dataset and benchmark creation, and expanding the application of *CL4Code* to a wider range of use cases.

Datasets and Benchmarks. Our findings in Section 4.1 highlight a substantial disparity in the performance of a PLM between ID and OOD data. Our results, along with a previous work [72], indicate that evaluating PLMs on ID data often leads to inflated metrics and results in overly optimistic conclusions in terms of the performance. Therefore, it is crucial to develop OOD datasets for code in order to evaluate the real-world generalizability of PLMs, as previously emphasized [69, 72]. Moreover, aligning dataset designs with continual learning scenarios offers the potential to evaluate the PLM’s ability to adapt to changing environments, which is crucial for practical deployment.

Improving benchmarks for PLMs of code is another promising direction for future research. Benchmarks such as CodeXGlue [41] play a crucial role by providing standardized evaluations of models of code and enabling reproducible experimental results. However, as such researches progress at a rapid pace, widely used benchmarks often become outdated quickly. In particular, Kiela et al. [33] showed that benchmarks such as GLUE [60] in NLP saturate, meaning the milestones set by the benchmark are reached. Thus, continued efforts to enhance benchmarks in deep learning for code are vital in establishing concrete goals and driving research to enhance the performance of the models being evaluated. Recently, Yang et al. [69] proposed GLUE-X, a comprehensive benchmark consisting of 13 datasets to test PLMs on OOD data across eight NLP tasks. The benchmark includes OOD datasets that are distinct from those in the original GLUE benchmark. Developing OOD benchmarks for code similar to GLUE-X [69] would greatly contribute to the growth of research on OOD generalization for PLMs of code. One potential approach is to compile a new set of OOD datasets that are not included in the existing CodeXGlue benchmark, and use them to test PLMs of code. Furthermore, exploring the design of OOD scenarios specific to software changes, as demonstrated in the present study, can provide a valuable foundation for future code benchmark initiatives. Our dataset and methodology for extracting OOD samples for API evolution scenarios can serve as a starting point for these endeavors.

Continual Learning for Code. Our findings in Section 4.2 highlight the challenge of catastrophic forgetting that PLMs of code encounter in a continual fine-tuning scenario with OOD data. Our study serves as a starting point for exploring the adaptability of PLMs of code in a variety of continual learning scenarios. For instance, these scenarios can be based on domain adaptation, where PLMs must adapt to new kinds of data such as new, unseen programming languages or code repositories as discussed in prior studies [25, 30, 35]. Additionally, incorporating continual learning into a multi-task learning framework is highly relevant to software engineering, given the multitude of downstream tasks involved.

In Section 4.3, our results demonstrate the effectiveness of continual learning methods in mitigating catastrophic forgetting in PLMs of code. We chose to explore these widely used methods as a first step in the research on continual learning for code. In the

future, more sophisticated techniques from NLP, as discussed in Section 6.2, can be evaluated. Furthermore, the creation of continual learning methods specifically tailored to source code has the potential to further reduce catastrophic forgetting in PLMs of code.

Finally, we did not focus our study on large language models (LLMs) as it would require a tremendous amount of available computational resources to pre-train an LLM from scratch under our OOD scenario. Nonetheless, we foresee that continuously adapting LLMs to new emerging datasets and benchmarks constitutes an exciting avenue for future work. In this context, and as fully fine-tuning LLMs is computationally costly, we believe that combining continual learning with parameter-efficient fine-tuning (PEFT) techniques might be beneficial to further enhance the capabilities of LLMs. These PEFT techniques have already shown promising results in LLMs for code intelligence [12, 61, 64].

6 RELATED WORK

6.1 Out-Of-Distribution Generalization

Natural Language Processing. Recent studies have revealed that PLMs are susceptible to generating inaccurate predictions when encountering OOD data [27, 54]. In NLP, this issue can manifest itself in situations where the domain of the test data differs from the pre-training data [23]. One approach to addressing this problem is to fine-tune PLMs on domain-specific datasets using efficient transfer learning techniques. For example, [29, 51] demonstrated that such approaches help PLMs in learning domain-specific knowledge and improve their generalization to unseen domains. Additionally, new datasets and benchmarks allow for further research on PLM domain adaptation. For instance, Williams et al. [66] introduced the MultiNLI dataset, containing text data from a variety of domains for PLM domain adaptation. Conneau et al. [14] proposed a cross-lingual NLI dataset for evaluating the cross-lingual transferability of PLMs. Recently, Yang et al. [69] introduced GLUE-X, a benchmark for evaluating PLMs’ ability to generalize to OOD data.

Deep Learning for Code. The study of OOD generalization of PLMs of code is an emerging research area. Assessing their generalizability and designing efficient techniques to improve their robustness to OOD scenarios is essential for the practical usability of PLMs of code [72]. Previous work in this field has focused on designing OOD datasets that simulate specific distribution shifts of program data. Koh et al. [35] presented PY150-Wilds, a Python dataset in which the test data consists of code repositories not appearing in the training data. The authors demonstrated performance gaps between the model on ID and OOD data. However, it is important to note that while the design choice is sound, it may not reflect strong OOD phenomena as the distribution of code tokens across different repositories may still be highly similar. More recently, Hu et al. [30] proposed a benchmark to evaluate the performance of code models under different distribution shift scenarios, including programmer, time, or token distribution shifts. In their study, the authors found that PLMs such as CodeBERT were robust against distribution shifts. However, they demonstrated that on a simple classification task with small datasets. In addition, the authors did not control the pre-training data of the studied PLMs, which can result in important data leakage between the pre-training and OOD

test data. This problem of data leakage is critical as some of the test data may have been seen by the model during pre-training. Overall, this is a prime threat to the validity of the OOD scenario that may lead to obtaining inflated metrics on the OOD test data. Finally, Hajipour et al. [25] analyzed the performance of PLMs of code on a syntax-based, semantic-based and complexity-based OOD scenario and highlighted that the models exhibit poor generalizability when faced with OOD samples. However, it is important to point out that the OOD scenarios used in this study may be too artificial. For instance, in the syntax-based scenario, some language-specific tokens are masked at training to study how the model generalizes to unseen language tokens. Such a scenario is unrealistic as it does not reflect the nature of OOD data that a PLM of code is likely to encounter in the real world. Additionally, there is no practical motivation for masking specific tokens while training the model.

In this study, we propose an OOD dataset that accurately represents the dynamic nature of software codebases in the real world. Specifically, we focus on the scenario where a PLM must adapt to new, unseen APIs over time, a well-established problem in the literature [46, 48]. To ensure the validity of our experiments, we thoroughly control our PLM setup to prevent any data leakage between the pre-training, fine-tuning, and test data. This allows us to create an OOD generalization scenario that is as close to reality as possible, an aspect that has been overlooked in previous works.

6.2 Continual Learning for Pre-trained Language Models

Continual learning has been studied to adapt pre-trained language models based on the Transformer architecture [57] to new domains or tasks in NLP. For example, Cao et al. [9] proposed a method to continually learn from new classes of events in textual data to detect them without degradation of the accuracy over time. Douillard et al. [17] introduced DyTox, a method that utilizes an encoder-decoder transformer for multiple tasks by expanding the network with task-specific special tokens, allowing for continual learning of new tasks with a low computational and memory footprint. Ermis et al. [18] proposed a memory-efficient approach for transformers to continually learn new tasks by sharing information across tasks and expanding the network with task-specific modules. Similarly, Vladymyrov et al. [58] proposed the HyperTransformer architecture to continually learn new tasks by generating task-specific convolutional neural network weights in a few-shot learning setting and updating the task-specific weights to avoid catastrophic forgetting. Lastly, Jie et al. [32] leverage continual learning to avoid representational shifts in PLMs by proposing a new hierarchical fine-tuning method that prevents excessive changes in the representation spaces of the neural network in a continual fine-tuning setting.

Recent advances in NLP highlight the crucial need for PLMs to adapt to changing environments and maintain their performance on new data and tasks. In the field of software engineering, the application of continual learning to PLMs of code is essential for developing methods that enable the model to robustly adapt to new codebases and tasks over time. To the best of our knowledge, only a couple of prior studies utilized continual learning in the context of

code intelligence. Baudry et al. [5] demonstrate the benefits of leveraging continual learning to fix bugs when considering a continuous stream of code change with continuous integration development platforms. The scope of our study differ from this prior work in many aspects. First, contrary to this prior work, our study focuses on PLM architectures which broadens the potential applicability of our approach to a broader range of tasks. Secondly, we compare numerous continual learning techniques in our OOD scenario with PLMs, whereas this previous work only consider a continual learning scenario without leveraging continual learning techniques such as replay buffer or EWC. More recently, Gao et al. [21] made similar findings than ours by showing that PLMs suffer from catastrophic forgetting in continual learning scenarios and that replay-based approaches allow to effectively mitigate forgetting. We believe that these prior works and our study break new ground by introducing the first approaches on the utilization of continual learning for PLMs of code.

7 CONCLUSION AND FUTURE WORK

Our study exposes the limitations of pre-trained language models of code in handling out-of-distribution data in a continual fine-tuning scenario. Our results reveal that OOD data significantly decreases the PLMs' effectiveness in two API-related downstream tasks compared to ID data. Our findings indicate that classical transfer learning fails to adapt the PLMs to new, unseen APIs in this evolution scenario. Additionally, we observe instances of catastrophic forgetting, prompting us to explore methods that address this issue. In our final experiments, we demonstrate that replay-based and regularization-based continual learning techniques can effectively mitigate catastrophic forgetting while retaining or enhancing the performance of the PLMs in both downstream tasks. In future work, we intend to explore more OOD scenarios to further evaluate the generalizability of PLMs of code and develop relevant OOD generalization benchmarks for code. Additionally, we plan to implement more advanced continual learning methods tailored to source code to enhance the adaptability of PLMs of code. Finally, we aim to investigate OOD detection methods to automatically identify OOD data in PLMs, thereby improving their performance.

DATA AVAILABILITY

We publicly release all the code, data and models to reproduce the experiments of our study. The following repository contains instructions on how to acquire the data and pre-train, fine-tune and test the PLMs: <https://github.com/martin-wey/cl-code-apis>

REFERENCES

- [1] Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified Pre-training for Program Understanding and Generation. In *Proceedings of the 2021 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] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*. 139–154.
- [3] Miltiadis Allamanis. 2019. The Adverse Effects of Code Duplication in Machine Learning Models of Code. In *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Athens, Greece) (Onward! 2019)*. Association for Computing Machinery, New York, NY, USA, 143–153. <https://doi.org/10.1145/3359591.3359735>

- [4] Gareth Ari Aye, Seohyun Kim, and Hongyu Li. 2021. Learning Autocompletion from Real-World Datasets. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. 131–139. <https://doi.org/10.1109/ICSE-SEIP52600.2021.00022>
- [5] Benoit Baudry, Zimin Chen, Khashayar Etemadi, Han Fu, Davide Ginelli, Steve Kommrusch, Matias Martinez, Martin Monperrus, Javier Ron, He Ye, and Zhongxing Yu. 2021. A Software-Repair Robot Based on Continual Learning. *IEEE Software* 38, 4 (2021), 28–35. <https://doi.org/10.1109/MS.2021.3070743>
- [6] Chaitanya Baweja, Ben Glocker, and Konstantinos Kamnitsas. 2018. Towards continual learning in medical imaging. *arXiv preprint arXiv:1811.02496* (2018).
- [7] Magdalena Biesialska, Katarzyna Biesialska, and Marta R. Costa-jussà. 2020. Continual Lifelong Learning in Natural Language Processing: A Survey. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 6523–6541. <https://doi.org/10.18653/v1/2020.coling-main.574>
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
- [9] Pengfei Cao, Yubo Chen, Jun Zhao, and Taifeng Wang. 2020. Incremental Event Detection via Knowledge Consolidation Networks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 707–717. <https://doi.org/10.18653/v1/2020.emnlp-main.52>
- [10] Arslan Chaudhry, Puneet K Dokania, Thalayisingam Ajanthan, and Philip HS Torr. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 532–547.
- [11] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- [12] YunSeok Choi and Jee-Hyong Lee. 2023. CodePrompt: Task-Agnostic Prefix Tuning for Program and Language Generation. In *Findings of the Association for Computational Linguistics: ACL 2023*. 5282–5297.
- [13] Matteo Cinielli, Nathan Cooper, Luca Pascarella, Denys Poshyvanyk, Massimiliano Di Penta, and Gabriele Bavota. 2021. An Empirical Study on the Usage of BERT Models for Code Completion. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. 108–119. <https://doi.org/10.1109/MSR52588.2021.00024>
- [14] Alexis Conneau, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053* (2018).
- [15] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. 2019. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv preprint arXiv:1909.08383* 2, 6 (2019), 2.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [17] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. 2022. Dyttox: Transformers for continual learning with dynamic token expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9285–9295.
- [18] Beyza Ermis, Giovanni Zappella, Martin Wistuba, Aditya Rawal, and Cedric Archambeau. 2022. Memory Efficient Continual Learning with Transformers. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 10629–10642. https://proceedings.neurips.cc/paper_files/paper/2022/file/4522de4178b9db36b49aa26fad537cf-Paper-Conference.pdf
- [19] 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*. Association for Computational Linguistics, Online, 1536–1547. <https://doi.org/10.18653/v1/2020.findings-emnlp.139>
- [20] Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* 3, 4 (1999), 128–135.
- [21] Shuzheng Gao, Hongyu Zhang, Cuiyun Gao, and Chaozheng Wang. 2023. Keeping Pace with Ever-Increasing Data: Towards Continual Learning of Code Intelligence Models. *arXiv preprint arXiv:2302.03482* (2023).
- [22] Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. UniXcoder: Unified Cross-Modal Pre-training for Code Representation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 7212–7225. <https://doi.org/10.18653/v1/2022.acl-long.499>
- [23] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 8342–8360. <https://doi.org/10.18653/v1/2020.acl-main.740>
- [24] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. 2020. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences* 24, 12 (2020), 1028–1040. <https://doi.org/10.1016/j.tics.2020.09.004>
- [25] Hossein Hajipour, Ning Yu, Cristian-Alexandru Staicu, and Mario Fritz. 2022. SimSCOOD: Systematic Analysis of Out-of-Distribution Behavior of Source Code Models. *arXiv preprint arXiv:2210.04802* (2022).
- [26] Vincent J. Hellendoorn, Sebastian Proksch, Harald C. Gall, and Alberto Bacchelli. 2019. When Code Completion Fails: A Case Study on Real-World Completions. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 960–970. <https://doi.org/10.1109/ICSE.2019.00101>
- [27] Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. 2020. Pretrained Transformers Improve Out-of-Distribution Robustness. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 2744–2751. <https://doi.org/10.18653/v1/2020.acl-main.244>
- [28] Thong Hoang, Hong Jin Kang, David Lo, and Julia Lawall. 2020. CC2Vec: Distributed Representations of Code Changes. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (Seoul, South Korea) (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 518–529. <https://doi.org/10.1145/3377811.3380361>
- [29] 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>
- [30] Qiang Hu, Yuejun Guo, Xiaofei Xie, Maxime Cordy, Lei Ma, Mike Papadakis, and Yves Le Traou. 2022. CodeS: A Distribution Shift Benchmark Dataset for Source Code Learning. *arXiv preprint arXiv:2206.05480* (2022).
- [31] Xing Hu, Ge Li, Xin Xia, David Lo, Shuai Lu, and Zhi Jin. 2018. Summarizing Source Code with Transferred API Knowledge. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (Stockholm, Sweden) (IJCAI'18)*. AAAI Press, 2269–2275.
- [32] S. Jie, Z. Deng, and Z. Li. 2022. Alleviating Representational Shift for Continual Fine-tuning. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE Computer Society, Los Alamitos, CA, USA, 3809–3818. <https://doi.org/10.1109/CVPRW56347.2022.00426>
- [33] Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, et al. 2021. Dynabench: Rethinking benchmarking in NLP. *arXiv preprint arXiv:2104.14337* (2021).
- [34] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114, 13 (2017), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
- [35] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. 2021. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*. PMLR, 5637–5664.
- [36] Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* 40, 12 (2017), 2935–2947. <https://doi.org/10.1109/TPAMI.2017.2773081>
- [37] 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. *arXiv preprint arXiv:1907.11692* (2019).
- [38] Vincenzo Lomonaco and Davide Maltoni. 2017. CORE50: a New Dataset and Benchmark for Continuous Object Recognition. In *Proceedings of the 1st Annual Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 78)*, Sergey Levine, Vincent Vanhoucke, and Ken Goldberg (Eds.). PMLR, 17–26. <https://proceedings.mlr.press/v78/lomonaco17a.html>
- [39] Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L. Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Gido van de Ven, Martin Mundt, Qi She, Keiland Cooper, Jeremy Forest, Eden Belouadah, Simone Calderara, German I. Parisi, Fabio Cuzzolin, Andreas Tolias, Simone Scardapane, Luca Antiga, Subutai Ahmad, Adrian Popescu, Christopher Kanan, Joost van de Weijer, Tinne Tuytelaars, Davide Bacciu, and Davide Maltoni. 2021. Avalanche: an End-to-End Library for Continual Learning. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (2nd Continual Learning in Computer Vision Workshop)*.
- [40] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. 2018. Learning under concept drift: A review. *IEEE Transactions on Knowledge and*

- Data Engineering* 31, 12 (2018), 2346–2363.
- [41] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, et al. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- [42] Fei Lv, Hongyu Zhang, Jian-guang Lou, Shaowei Wang, Dongmei Zhang, and Jianjun Zhao. 2015. CodeHow: Effective Code Search Based on API Understanding and Extended Boolean Model (E). In *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 260–270. <https://doi.org/10.1109/ASE.2015.42>
- [43] Michael McCloskey and Neal J. Cohen. 1989. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *Psychology of Learning and Motivation*, Vol. 24. Academic Press, 109–165. [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8)
- [44] Kawser Wazed Nafi, Tonny Shekha Kar, Banani Roy, Chanchal K. Roy, and Kevin A. Schneider. 2019. CLCDSA: Cross Language Code Clone Detection using Syntactical Features and API Documentation. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 1026–1037. <https://doi.org/10.1109/ASE.2019.00099>
- [45] Tung Thanh Nguyen, Hoan Anh Nguyen, Nam H. Pham, Jafar M. Al-Kofahi, and Tien N. Nguyen. 2009. Graph-Based Mining of Multiple Object Usage Patterns. In *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering (Amsterdam, The Netherlands) (ESEC/FSE '09)*. Association for Computing Machinery, New York, NY, USA, 383–392. <https://doi.org/10.1145/1595696.1595767>
- [46] Marius Nita and David Notkin. 2010. Using twinning to adapt programs to alternative APIs. In *2010 ACM/IEEE 32nd International Conference on Software Engineering*. Vol. 1. IEEE, 205–214.
- [47] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks* 113 (2019), 54–71. <https://doi.org/10.1016/j.neunet.2019.01.012>
- [48] Sebastian Proksch, Sven Amann, Sarah Nadi, and Mira Mezini. 2016. Evaluating the evaluations of code recommender systems: a reality check. In *2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 111–121.
- [49] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [50] Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. Codebleu: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297* (2020).
- [51] Sebastian Ruder, Matthew E. Peters, Swabha Swayamdipta, and Thomas Wolf. 2019. Transfer Learning in Natural Language Processing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*. Association for Computational Linguistics, Minneapolis, Minnesota, 15–18. <https://doi.org/10.18653/v1/N19-5004>
- [52] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671* (2016).
- [53] Zheyang Shen, Jiashuo Liu, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. 2021. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624* (2021).
- [54] Yuge Shi, Imant Daunhawer, Julia E Vogt, Philip Torr, and Amartya Sanyal. 2022. How robust are pre-trained models to distribution shift?. In *ICML 2022: Workshop on Spurious Correlations, Invariance and Stability*. <https://openreview.net/forum?id=zKDcZBVVEWm>
- [55] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. 2020. Few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12183–12192.
- [56] Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. Overcoming Catastrophic Forgetting During Domain Adaptation of Neural Machine Translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 2062–2068. <https://doi.org/10.18653/v1/N19-1209>
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
- [58] Max Vladymyrov, Andrey Zhmoginov, and Mark Sandler. 2023. Continual Few-Shot Learning Using HyperTransformers. *arXiv preprint arXiv:2301.04584* (2023).
- [59] Li Wan, Matthew Zeiler, Sixin Zhang, Yann LeCun, and Rob Fergus. 2013. Regularization of Neural Networks Using Dropconnect. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28 (Atlanta, GA, USA) (ICML '13)*. JMLR.org, III–1058–III–1066.
- [60] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461* (2018).
- [61] Chaozheng Wang, Yuanhang Yang, Cuiyun Gao, Yun Peng, Hongyu Zhang, and Michael R Lyu. 2022. No more fine-tuning? an experimental evaluation of prompt tuning in code intelligence. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 382–394.
- [62] 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 *Proceedings of the 2021 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>
- [63] Cody Watson, Nathan Cooper, David Nader Palacio, Kevin Moran, and Denys Poshyvanyk. 2022. A systematic literature review on the use of deep learning in software engineering research. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 31, 2 (2022), 1–58. <https://doi.org/10.1145/3485275>
- [64] Martin Weysow, Xin Zhou, Kisub Kim, David Lo, and Houari Sahraoui. 2023. Exploring Parameter-Efficient Fine-Tuning Techniques for Code Generation with Large Language Models. [arXiv:2308.10462](https://arxiv.org/abs/2308.10462) [cs.SE]
- [65] Gerhard Widmer and Miroslav Kubat. 1996. Learning in the presence of concept drift and hidden contexts. *Machine learning* 23, 1 (1996), 69–101. <https://doi.org/10.1023/A:1018046501280>
- [66] Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426* (2017).
- [67] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771* (2019).
- [68] Frank F Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. A systematic evaluation of large language models of code. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*. 1–10.
- [69] Linyi Yang, Shuibai Zhang, Libo Qin, Yafu Li, Yidong Wang, Hanmeng Liu, Jindong Wang, Xing Xie, and Yue Zhang. 2022. GLUE-X: Evaluating Natural Language Understanding Models from an Out-of-distribution Generalization Perspective. *arXiv preprint arXiv:2211.08073* (2022).
- [70] Zhengran Zeng, Hanzhuo Tan, Haotian Zhang, Jing Li, Yuqun Zhang, and Lingming Zhang. 2022. An extensive study on pre-trained models for program understanding and generation. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*. 39–51. <https://doi.org/10.1145/3533767.3534390>
- [71] Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*. PMLR, 3987–3995.
- [72] Xin Zhou, DongGyun Han, and David Lo. 2021. Assessing generalizability of CodeBERT. In *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 425–436.

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