



MergeRepair: An Exploratory Study on Merging Task-Specific Adapters in Code LLMs for Automated Program Repair

Meghdad Dehghan¹ · Jie JW Wu² · Fatemeh H. Fard³ · Ali Ouni⁴

Received: 30 May 2025 / Accepted: 20 July 2026
© The Author(s) 2026

Abstract

Large Language Models (LLMs) have shown high capabilities in several software development-related tasks such as program repair, documentation, code refactoring, debugging, and testing. However, training these models requires massive amount of data and significant computational resources. Adapters are specialized, small modules designed for parameter efficient fine-tuning of LLMs for specific tasks, domains, or applications without requiring extensive retraining of the entire model. These adapters offer a more efficient way to customize LLMs for particular needs, leveraging the pre-existing capabilities of the large model. Model (and adapter) merging have emerged as a technique to develop one model capable of multiple tasks, with minimal or no training required. Although model and adapter merging has shown promising performance in domains such as natural language processing and computer vision, its applicability to software engineering tasks remains underexplored. In this paper, we investigate the effectiveness of merged adapters within the context of software engineering, with a particular focus on the Automated Program Repair (APR) task, through our approach, **MergeRepair**. In particular, we merge multiple task-specific adapters using three different merging methods, including weight-averaging, TIES-Merging, and DARE-TIES, and evaluate the performance of the merged adapter on the APR task. We introduce a **continual merging** approach, a novel method in which we sequentially merge the task-specific adapters where the order and weight of the merged adapters play a significant role. We further compare the performance of our approach with a baseline method consisting of equal-weight merging applied on parameters of different adapters, where all adapters are of equal importance. Through continual merging, we explore the capability of merged adapters and the effect of task order, as it occurs in real-world software projects. To evaluate our approach, we consider two LLMs, StarCoder2 and Granite models on an APR benchmark, HumanEvalFix, with various combinations of the merged adapters trained on a subset of the CommitPackFT dataset that contains four tasks and APR in Python. The results show that merging task-specific adapters can enhance the performance of models on the APR task, *without additional training*. In particular, our approach achieves up to 2.38% performance improvement over APR task in terms of pass@1 score when merging APR with Improvement and Misc task-adapters on top of the StarCoder2 model. Performance rises to 4.01% for pass@10 score with the

Communicated by: Shinpei Hayashi, Shane McIntosh.

Extended author information available on the last page of the article

same model and merged adapters. This improvement could be significantly influenced by the base model and the specific task adapters merged, rather than only the number of adapters. The proposed continual merging approach can be effective if the order of merging is efficiently determined. Results show that performance improves when the most effective task-adapter is added last in a continual merging process. Moreover, we can achieve on-par or even better performance when APR adapter is not used in the merged model, emphasizing the generalizeability of the approach. These findings suggest that leveraging knowledge from other code-related tasks through adapter merging can improve performance on new tasks like APR, even with limited data.

Keywords Automated program repair · Model merging · Code LLMs · Adapters

1 Introduction

Significant Software Engineering (SE) efforts have been devoted to automate code-related tasks and facilitate developers' tasks using deep learning models. Bug prediction (Pandey et al. 2020), bug fixing (Chen et al. 2019; Lutellier et al. 2020), code generation (Chen et al. 2021; Nijkamp et al. 2023), comment generation (Hu et al. 2020), and commit message generation (Liu et al. 2020) are among these tasks. In practice, developers often need to train, test and validate different models for each of these tasks separately along with data collection and pre-processing, feature engineering, hyper-parameter tuning, and so on, to ensure efficient models are developed. However, repeating the process for each new task can be time-consuming and cumbersome. To train a new model that can perform well on a downstream task, there are several requirements, including an (often domain-specific) dataset, and sufficient computational resources. Prior research has shown that these requirements are not trivial to fulfill (Liu et al. 2024). First, for each new task, a new model should be trained and deployed, which is costly. Additionally, by following this process, practitioners will end up with several task-specific models, which are unable to leverage the strengths of other models, while they could benefit from related tasks to improve the performance within or outside their domain (Goddard et al. 2024; Stoica et al. 2024).

The main question that we intend to investigate in this study is “*Can we re-use previously trained code-related adapters for a new SE task by merging them?*” Adapters are small modules used in Parameter-Efficient Fine-Tuning (PEFT) (Houlsby et al. 2019a) of Large Language Models (LLM) and have results on par with or better than fully fine tuning language models (Ding et al. 2023; Weyssow et al. 2023a). The idea behind our study is rooted in the foundations of transfer learning (Torrey and Shavlik 2010), where the learned knowledge of a language model can be adapted to new tasks, domains, or languages (Raffel et al. 2020). The advantages of transfer learning in software engineering and code intelligence have been shown by several recent works adopting LLMs (Silva et al. 2023; Weyssow et al. 2023a). Although fine-tuning is one of the main approaches for using language models, there are concerns about using several task-specific adapters or models. The issue of catastrophic forgetting (Weyssow et al. 2023b) is another concern when fine-tuning a language model. Additionally, fully fine-tuning LLMs (Howard and Ruder 2018) becomes inefficient as multiple instances of the model for different tasks should be trained and deployed.

To address the above-mentioned issues, recent research showed the efficiency of merging models. Merging models enhances their performance compared to task-specific models (Goddard et al. 2024; Stoica et al. 2024). Goddard et al. (2024) showed that merging models combines the parameters of individual task-specific models into a single model in order to leverage the knowledge of other models. Note that merging models is intrinsically different from Multi-Task Learning (MTL) (Zhang and Yang 2018; Liu et al. 2015). In MTL, one model is jointly trained on two or more tasks; while in merging models, there exist multiple models, each of them trained separately on distinct tasks. Then, all of these models are merged together. Though there are multiple approaches for model merging, we choose the techniques that merge models *without additional training* (Yang et al. 2024).

Similar techniques are developed for merging adapters. There are studies that merge multiple models (Matena and Raffel 2022; Ainsworth et al. 2022; Yu et al. 2023) or adapters in natural language processing (NLP) domain (Chronopoulou et al. 2023; Nguyen and Le 2024). Adapters were also used for Multi-Task Learning (Pfeiffer et al. 2021). Others adopted the idea of mixture-of-experts to inject a set of expert adapter layers in the transformer-based models in order to train multi-task models (Wang et al. 2022). However, the research on merging adapters is scarce, and there is a research gap in investigating the merging ability of adapters for code-related tasks. This is specially important with the advent of Code LLMs—LLMs that are pre-trained specifically to understand and generate code in various programming languages. Though these models have shown promising results for many SE tasks, their computational cost is not negligible for companies and researchers. Thus, it is beneficial to re-use the trained models/adapters for new tasks by merging them, *without additional training*. There is no prior research that have investigated whether current approaches affect code-related tasks in the same way as in NLP and whether they would improve the performance of each task in the merged model.

The main goal of this study is to *explore the performance of merged task-specific adapters and continual merging, in the context of Automated Program Repair (APR) task*. We investigated the effect of merging adapters for different SE tasks, using our framework, **MergeRepair**. First, we examined three merging techniques from the existing literature, being weight-space averaging (Wortsman et al. 2022), TIES-Merging (Yadav et al. 2024) and DARE-TIES, based on DARE (Yu et al. 2023), in which all adapters are of *equal* importance in the merged model. We studied the effect of merging on APR, and the generalizability of merging adapters to APR, i.e., when the merged model is tested on APR without being trained on it. Second, we explored our proposed merging paradigm, **continual merging** using all three merging techniques (i.e., weight-space averaging, TIES-Merging, and DARE-TIES). In continual merging, the *order and weights* of the adapters play an important factor, and the merged adapters do not have equal importance. Our proposed approach is motivated by the practical considerations of merged models, when new models are often introduced sequentially. Additionally, not all tasks could benefit the new target task equally. The continual merging addresses these two points by considering different weights based on the order of the tasks in the merged model.

To evaluate our approach, we utilize the Low-Rank Adaptation (LoRA) (Hu et al. 2022) as adapter modules and train one instance of LoRA per task. LoRA optimizes low-rank decomposition of the weight matrices and, in this way, can achieve comparable results to fully fine-tuning (Ding et al. 2023; Weyssow et al. 2023a). In our framework, **MergeRe-**

pair, we explore the merging capability of adapters injected in Code LLMs to improve the performance of a selected task, automatic program repair. APR is an active research area aiming to reduce the manual effort for software developers (Liu et al. 2021), and facilitates software development and maintenance (Goues et al. 2012). Specifically, APR automates the process of generating the fixed code or patch, for a given code that contains bugs. We chose the broader categories of the tasks from CommitPackFT dataset (Muennighoff et al. 2023), being Development, Test & QA, Improvement, and Misc (includes tasks like configuration, formatting, dependencies, and documentation). CommitPackFT dataset (Muennighoff et al. 2023) is chosen as it is a widely used dataset for code related tasks and code repair (Lozhkov et al. 2024; Muennighoff et al. 2023; Zhuo et al. 2024) for our study, as merging tasks to APR (i.e., Bug category of the CommitPackFT dataset). The dataset contains 59K samples, with $\sim 19\%$ belonging to APR.

Findings Our results show that the base models and the constituent tasks of the merged adapter play an important role in improving the performance of APR in a merged model. Additionally, the performance of the merged adapter depends more on the task-adapter being used in the merging than the number of task-adapters being merged. Furthermore, we observe that merged adapters originating from non-APR tasks can still perform well on the APR benchmark, provided their constituent adapters achieve strong APR scores. In some cases, the merged adapter even outperforms the APR-specific adapter, demonstrating the generalizability of the merging approach. Finally, in the continual merging approach, the order of merging task-specific adapters affects the performance of the merged adapters significantly. For most of the merged task-specific adapters, we observed an improved performance when using continual merging, considering the best order of the adapter, compared to their equal-weight averaging.

The contributions of this study are as follows:

- We are the first to study merging adapters for automatic program repair task.
- We empirically study the effect of different equal-weight averaging merging techniques on automatic program repair task, using our framework, **MergeRepair**.
- We propose continual merging, a new technique that considers different weights for the adapters in the merged model.
- Our code and replication package are made public¹ to support open data and open science principles.

Our study offers actionable insights for both practitioners and researchers interested in model reuse through merging techniques across various software engineering tasks. Since the merging methods explored in this work do not require training on datasets, they can support data confidentiality by enabling model sharing without exposing sensitive data. Additionally, our findings highlight the role of different tasks in influencing the effectiveness of model merging for automated program repair, potentially opening new avenues for future research.

The rest of the paper is structured as follows. Section 2 describes the background and related work of our research. Section 3 explains the design of our empirical study. Section

¹Replication package: <https://github.com/mqddd/mergerepair>

4 summarizes the results for RQs. Section 5 includes more analysis and discussions on the results. Implications of study and threats to validity are explained in sections 6 and 7, respectively. Finally, Section 8 concludes this work.

2 Related works

2.1 Automated program repair

Early research in APR includes searched-based methods which are studied in several works (Weimer et al. 2010, 2009; Motwani and Brun 2020; Wong et al. 2021; Jobstmann et al. 2005). These studies search for the most probable places in the code that are prone to be buggy and then generate candidate patches to fix the buggy codes. To localize the buggy part, various techniques are utilized. In (Jobstmann et al. 2005), the authors treat the program repair task as a game and propose a heuristic method to search in the search space. More recent techniques such as SimFix (Motwani and Brun 2020) and VarFix (Wong et al. 2021) use code similarity metrics to find the correct patch among a set of candidate patches. More specifically, SimFix (Motwani and Brun 2020) uses the Abstract Syntax Tree (AST) of the similar code snippets and existing candidate patches and searches in the intersection of these two sets to find the correct patch. VarFix (Wong et al. 2021) uses the variational execution technique in which a program is executed multiple times that allows systematically exploration in large search spaces.

Constraint-based and template-based methods are other approaches for APR that are introduced following search-based methods. In constraint-based methods, the focus is on generating patch candidates based on a set of rules or constraints. Such techniques make the search space more limited compared to other methods (Wei et al. 2010; Nguyen et al. 2013; Xuan et al. 2016). Template-based methods are introduced to alleviate the issues with the previous methods. Such tools utilize a predefined program template for generating fixed patches. For example, in FixMiner (Koyuncu et al. 2020), authors proposed an approach to mine fix patterns leveraging the changes within patches that could be used in patch generation systems. Later works on automatic program repair train deep neural net (Lutellier et al. 2020; Mesbah et al. 2019) and use pre-trained language models (Fu et al. 2022; Xia et al. 2023). By the emergence of LLMs, the research towards adapting these models to downstream SE tasks has shifted to instruction-tuning them to be aligned with the target domain of interest (Fu et al. 2022; Xia et al. 2023).

In this research, to develop the task-specific models, we train adapters using instruction-tuning datasets and evaluate the models, containing the merged adapter, on the APR task. Although language models and PEFT methods are studied for APR in previous works (Silva et al. 2023), there is no current research that leverages the combination of similar task-specific adapters using merging methods in order to study the performance of the APR task.

2.2 Adapters

Adapters are light-weight modules that are inserted between the Transformer layers and used as a technique for parameter-efficient fine-tuning (Houlsby et al. 2019b). Adapters have emerged as a technique to optimize memory usage by reducing the number of trainable

parameters during fine-tuning (Houlsby et al. 2019a). They have shown promising results compared to fully fine-tuned models while efficiently adapting large language models to downstream tasks. Subsequently, several adapter architectures have been proposed, including LoRA (Hu et al. 2022) and IA3 (Liu et al. 2022), each targeting specific layers and parameters of the transformer-based models. There are empirical studies on the utilization of adapters for SE tasks (Liu et al. 2024; Wang et al. 2023a). Notably, recent studies have adopted adapters as a primary alternative method for efficiently fine-tuning large language models in specialized domains such as program repair (Silva et al. 2023). Apart from APR and SE tasks, adapters and PEFT methods have been utilized primarily in other domains such as natural language processing (Houlsby et al. 2019a; Poth et al. 2023), computer vision (Chen et al. 2024; Sung et al. 2022; Xin et al. 2024), and speech recognition (Li et al. 2023; Otake et al. 2023). Though adapters and LORA have been used in several software engineering studies (Zou et al. 2025; Haque et al. 2025; Männistö et al. 2025), merging them for APR is not studied.

In this study, we opt for LoRA as the adapter for fine-tuning the models across all target tasks. LoRA stands out as one of the most prevalent and widely used methods in the research community (Ding et al. 2023; Weyssow et al. 2023a). We then apply the merging techniques on the trained LORA adapters.

2.3 Model merging

Merging models involves combining the parameters of two or more specialized models to create a unified multi-task model (Matena and Raffel 2022). Though multi-task model is one of the main motivations of model merging, this technique has been used for model unlearning, knowledge transfer among different languages, detoxification of LLMs, adversarial learning, federated learning, and multi-modal and cross-modal fusions as well (Yang et al. 2024). Zhang et al. (2023) proposed to compose different PEFT modules with arithmetic operations without retraining and demonstrated the benefits of such new PEFT in several experiments. Several approaches are proposed to enhance model merging, including EMR-merging (Huang et al. 2024), LoraHub (Huang et al. 2023), mixture of LORA (Wu et al. 2024), AdaMerging (Yang et al. 2023), AdaMergX (Zhao et al. 2024), and LoraSoups (Prabhakar et al. 2024). These works merge models or PEFT with or without training, each trying to address some limitations of the existing works.

Recent work has demonstrated that merging Llama2-7b-chat, a general-purpose chat model, with Meditron-7b, specialized for the medical domain, resulted in a merged model that outperformed its constituent models across both general and medical benchmarks (Goddard et al. 2024). The increasing number of merged models on the Open LLM leaderboard (Beeching et al. 2023) further proves the success of applying this approach to various benchmarks. Several studies have proposed different merging techniques, such as weight-space averaging (Wortsman et al. 2022), Fisher-Weighted averaging (Matena and Raffel 2022), Git Re-Basin (Ainsworth et al. 2022), TIES-Merging (Yadav et al. 2024), and DARE (Yu et al. 2023), to efficiently merge models trained on distinct domains. Considering the trend of scaling models to billions of parameters and the cost of fully fine-tuning these models, there is an increasing need to explore the feasibility of merging multiple adapters within a single language model when adapters are used to fine-tune the models for the SE domain. Works such as Adaptersoup (Chronopoulou et al. 2023) have investigated this approach for

both in-domain and out-of-domain evaluation of mixed adapters. Another study has conducted a comprehensive analysis of the effectiveness of different adapter architectures for the in-domain generalizability of merged adapters (Nguyen and Le 2024).

In this work, we employ three merging methods, i.e., weight-space merging (Wortsman et al. 2022), TIES-Merging (Yadav et al. 2024), and DARE-TIES (Yu et al. 2023), to merge multiple task-specific adapters. Specifically, we train multiple LoRA instances on various SE tasks and assess the performance of the merged LoRA in equal-weight merging and continual merging scenarios for all three merging methods. Though there are several model merging studies that merge models or PEFT approaches with or without training, there is no work that investigates the applicability of this technique in software engineering and specifically for automatic program repair.

3 Study design

3.1 Research questions

The main goal of this study is to investigate the performance of merged task-specific adapters on the APR task. Specifically, we use two main scenarios for merging adapters 1) equal-weight merging (RQ1 and RQ2) and 2) continual merging (RQ3), as shown in Fig. 1. The explanation of each scenario is provided in Section 3.3. We design our experiments to answer the following research questions:

RQ1: Does merging APR adapter with other task-specific adapters improve the performance of APR? We are interested to investigate whether the performance of the APR task will be degraded or improved if we merge APR adapter with non-APR adapters.

Prior studies have shown that merging adapters can improve the performance of single tasks in other domains (Goddard et al. 2024; Yang et al. 2024). In this RQ, we intend to investigate the potential improvement of merged adapters on the APR task, when APR is merged with four tasks of Development, Misc, Test & QA, and Improvement. The alternative way to achieve better performance having an already fine-tuned model, is to continue the fine-tuning process with additional data and computation. The potential improvement using merging adapters is essentially interesting because the merging process of various




Fig. 1 An overview of MergeRepair

adapters is computationally cheaper than further fine-tuning of models. To study this RQ, we start by merging APR with only one task and evaluate it on the APR benchmark. Thereafter, we continue merging APR adapters with two, three, and all four tasks and assess the performance of the merged adapters on APR. In such a manner, we consider all subsets of four adapters merged with APR adapter.

RQ2: Can merged adapters of other tasks generalize to APR? Our main interest for this RQ is to understand if the performance of other task-specific adapters could be maintained for the APR task when merged together.

Training a Code LLM on a new task is costly, and one of the goals of merging adapters is to expand the model to new tasks (i.e., out of domain data) without additional training (Godard et al. 2024). Hence, we aim to analyze the generalizability of merged adapters to out-of-domain data. Particularly, we merge non-APR adapters together and evaluate them on the APR benchmark. For merging task-specific adapters, similar to RQ1, we start with adapters trained on a single task and experiment with merging up to four adapters. Each time, we test the merged model on the APR dataset.

RQ3: How does continual merging of other task-specific adapters with the APR adapter influence the performance of APR? In real-world software projects, accessing several task-adapters at a time is less applicable; instead, new datasets and tasks will be available as the project evolves. In such situations, the ability to adapt the merged adapter to new data and tasks becomes crucial.

To this end, we investigate the continual learning capacity of merged adapters in this RQ, using our proposed paradigm, *continual merging*. Furthermore, the order in which different task-specific adapters will be merged is important as it determines their influence in the merged adapter (Section 3.3.1 shows the weight contribution or influence of different task-adapters for this scenario). We consider all permutations of the task-adapters to evaluate their order on APR performance.

3.2 Datasets

We utilize the CommitPackFT dataset released by the OctoPack study (Muennighoff et al. 2023). This dataset is a refined version of the full CommitPack dataset, filtered by its publishers to make it suitable for instruction-tuning Code LLMs (Muennighoff et al. 2023). CommitPackFT contains 277 programming languages and is 2GB of memory (Muennighoff et al. 2023). These two datasets contain code samples from GitHub repositories, including the code before and after the commit change. In particular, the dataset includes three main fields: 1) commit messages, 2) old contents (the file content before the commit), and 3) new contents (the file content after the commit). The commit messages are the instructions used when instruction-tuning the models. Generally, these commits enhance the previous version of the code. A portion of the commits are related to bug-fixing commits tailored for the program repair task.

The Python split of the CommitPackFT dataset has been classified into five tasks by the publishers of the dataset. This classification shows the task information for each commit, which is done by using 1-shot prompting with the GPT-4 model. The task classification for other programming languages is not available. Therefore, our experiments are limited to the Python split of the dataset. We will conduct all the experiments on the Python split of the dataset, which contains 59, 113 samples. Table 1 shows the available tasks in the Python

Table 1 Task distribution of the CommitPackFT dataset

Task	# Records
Development	15,656 (26.49%)
Bug Fixes (APR)	11,260 (19.02%)
Misc	11,617 (19.65%)
Test & QA	8,255 (13.96%)
Improvement	12,325 (20.85%)

dataset. This dataset is used to train the task-specific adapters. The tasks include Development, Bug Fixes, Misc, Test & QA, and Improvement, where each one except Misc contains a main task and a small portion of the dataset as a closely related task. Bug Fixes is used for APR and includes bug fixes. Development mainly contains new features and 0.88% of the whole data on user interfaces. Test & QA includes mainly testing and a small subset on logging/instrumentation (0.62% of the whole data). Improvement contains data mainly tagged as refactoring/code clean up, with small percentages labeled as formatting/linting (0.40% of the whole data) and performance improvement (0.64% of the whole data). The Misc subset contains a more variety of the tasks, including deprecation, build system/tooling, documentation, dependencies, configuration, and release management.

3.3 Experiment design

An overview of our approach, MergeRepair, is shown in Fig. 1. MergeRepair consists of three main phases: tuning phase, merging phase and the evaluation phase.

- **Tuning phase:** First, we train one instance of LoRA adapter for each task in the tuning phase.
- **Merging phase:** Then, we merge these task-specific adapters using two merging scenarios/paradigms: weight-space merging and continual merging. For the weight-space merging, we merge the parameters of all task-specific adapters with equivalent weights (i.e., influence) to investigate the potential improvement on APR (**RQ1**), generalizability to APR (**RQ2**), and the effect of continual merging (**RQ3**). This is done using three merging techniques. Subsequently, we merge non-APR adapters, again with equal weights.
- **Evaluation phase:** we evaluate the merged adapter on the APR benchmark to study the above-mentioned RQs.

For the continual merging scenario, we aim to simulate the continual learning capability of merged adapters in which we sequentially add new task-specific adapters to the previous merged adapter. In the following sections, we will explain the continual merging, merging methods, tasks, models, steps to answer each RQ, and evaluation metrics that we employ to study the above-mentioned RQs.

3.3.1 Continual merging

Let the parameters of the Code LLM be denoted by θ_M . We keep these parameters unchanged during instruction-tuning. For each target task, we have a separate LoRA adapter on top of the frozen model trained on that task. We refer to these trained adapters as task-specific

adapters. Formally, having n available tasks denoted by $T_i = \{T_1, T_2, \dots, T_n\}$, we will also have n adapters, each with distinct weight parameter values. We define I_{T_i} as the influence of adapter T_i in the merged adapter.

For continual merging, we maintain a single adapter as the merged adapter. Subsequently, we merge the parameters of the current merged adapter with one new task-specific adapter at a time. Therefore, starting from the first task, i.e., T_1 we will have $I_{T_1} = 1$ in the first step. In the second step, we add the second task-specific adapter, i.e., T_2 , resulting in $I_{T_1} = \frac{1}{2}$ and $I_{T_2} = \frac{1}{2}$ in the merged adapter. In the final step, n , the influence of all previously added adapters to the merged adapter will be similar to a geometric progression with $I_{T_n} = \frac{1}{2}$ for the last adapter and ratio of $\frac{1}{2}$. More formally, the final influence of task-specific adapters in the merged adapter, starting from T_1 and ending with T_n , would be acquired by

$$I_{T_i} = \frac{1}{2^{n+1-i}} \quad (1)$$

for $i = \{2, 3, \dots, n\}$ and $I_{T_i} = \frac{1}{2^{n-i}}$ for $i = 1$, since we consider the first adapter as the merged adapter in the first step.

Consequently, in this scenario, the influence of the parameters of the previous task-adapters in the merged adapter will decrease, as we reach the final (i.e., n -th) adapter. Intuitively, the merged adapter attends to the new task while retaining less knowledge about the previous tasks.

For the continual merging scenario, we consider different orders of task-specific adapters in this study. The order of adding task-specific adapters is important as it will lead to different weight values for the parameters of the merged adapter. Having n different adapters, the total number of their permutations would be equal to $n!$. Hence, to adhere to the computational resources available to us, for the continual merging scenario, we will merge APR adapter with up to three other task-specific adapters, totaling $4! = 24$ experiments required for each Code LLM.

3.3.2 Merging Methods

We experiment with three merging methods, all compatible with LoRA adapter that we aim to employ. Weight-space averaging is selected as the baseline merging method. The other two methods TIES-Merging (Yadav et al. 2024) and DARE-TIES based on DARE (Yu et al. 2023) propose more efficient merging techniques, which have been shown to perform better than weight-space averaging in prior research (Goddard et al. 2024).

Weight-space averaging This method is the most straightforward approach to merge multiple models as it averages the weight parameters of the specialized models in order to obtain the parameters of the merged model (Wortsman et al. 2022). For weight-space averaging, the formal influence of task-specific adapter T_i , is defined as

$$I_{T_i} = \frac{1}{n} \quad (2)$$

where $i = \{1, 2, 3, \dots, n\}$. This approach assigns equal importance to all adapters in the parameter space, creating a simple average of their weights.

TIES-Merging This method uses a three-step approach to merge two or more models (Yadav et al. 2024). Considering that all models involved in the merging process share the same architecture, TIES-Merging first aligns the models so that each parameter corresponds to its counterpart in the other models. It then removes parameters with smaller values or magnitudes compared to their corresponding parameters. This step aims to retain only the most influential parameters while reducing the impact of low-influence ones. The specialized fine-tuned parameters of a model are represented as task vectors. After obtaining trimmed task vectors for each model, a sign value of -1 or $+1$ is assigned to the remaining parameters by resolving sign conflicts. The sign is chosen to match the sign of the larger parameters in the trimmed task vector. Finally, parameters whose signs align with the sign values are averaged to form the final parameter vector of the merged model.

DARE-TIES DARE (Yu et al. 2023) uses a two-step procedure. Starting with a pre-trained LM with parameters θ_{PRE} , supervised fine-tuning produces an optimized model with parameters θ_{SFT} . For k tasks $t_1, t_2, \dots, t_k$, DARE defines delta values as the difference between the fine-tuned and pre-trained parameters, formulated as:

$$\delta^t = \theta_{SFT}^t - \theta_{PRE} \quad (3)$$

After deriving the δ values for each task, DARE applies its two-step procedure to obtain task-specific models. First, it drops a portion of the δ parameters by setting them to 0, using a random rate p . This is based on the claim that not all parameters changed during fine-tuning capture the downstream task's properties, so some remain redundant. Setting these to 0 still preserves the task-specific nature of the SFT model. After this dropping phase, DARE rescales the remaining parameters by a factor of $1/(1 - p)$. This can be expressed as:

$$\begin{aligned} m^t &\sim \text{Bernoulli}(p), \\ \tilde{\delta}^t &= (1 - m^t) \odot \delta^t, \\ \hat{\delta}^t &= \tilde{\delta}^t / (1 - p) \end{aligned}$$

Furthermore, to obtain the parameters of a DARE model, $\hat{\delta}_t$ is added to θ_{PRE} . The parameters of multiple DARE models can then be fused using any merging method, such as the previously mentioned weight-space averaging or TIES-Merging, to derive the parameters of the merged model. In this paper, we use DARE combined with TIES-Merging for our experiments, referred to as DARE-TIES, as TIES-Merging provides a more effective approach compared with weight-space averaging.

3.4 Tasks

Table 1 shows all the tasks, including Development, Misc, Test & QA, and Improvement, along with the number of their samples. The “Bug Fixes” portion of this data is associated with the APR task and forms 19.02% of the dataset. The Misc category includes code

samples from different domains, including configuration, dependencies, and documentation (Muennighoff et al. 2023). Development includes records related to new features added to the programs along with a small portion of the user interface. Improvement contains samples from refactoring/code cleanup and performance improvement. The Test & QA includes samples from testing and logging/instrumentation. For more details on the tasks refer to the OctoPack (Muennighoff et al. 2023) study. For all these tasks, we used instruction-tuning data to fine-tune the Code LLMs using LoRA. This training provides the task-specific adapters for our merging experiments. The pre-trained models, without instruction-tuning, will generate open-ended outputs that are not suitable for production use. This process aims to adapt the pre-trained models to generate human-aligned outputs given an instruction and an input/output pair of code samples (Ouyang et al. 2022a).

3.5 Models

We will use two state-of-the-art Code LLMs, StarCoder2 (Lozhkov et al. 2024) and Granite (Mishra et al. 2024), for our experiments. These models are chosen as they are more recent than popular Code LLMs like CodeLlama and Mistral. More importantly, StarCoder2 and Granite have the best performance for APR task at the time of this study (Lozhkov et al. 2024; Mishra et al. 2024). Both of these models have the same size, with three billion parameters each. Their checkpoints and source code are publicly available, allowing us to inject low-rank matrices of LoRA layers into the decoder blocks of these models and train only the added parameters. Another reason to choose StarCoder2 and Granite over other CodeLLMs is our restricted computational capacity. We ran experiments to fine-tune the two models using LoRA with a configuration that is reported by the models' authors to ensure the feasibility of our approach and its performance. This choice is not feasible for us with CodeLlama and Mistral.

starcoder2-3b StarCoder2 (Lozhkov et al. 2024) is a category of models developed and released by the open-source community as part of the BigCode project. These models are published in three sizes: 3b, 7b, and 15b. They are trained on the Stack v2 dataset, which is built on top of the Software Heritage (SWH) repositories. They also utilize a comprehensive set of publicly available data sources such as GitHub issues, StackOverflow, Jupyter and Kaggle notebooks. The architecture of these models is based on the StarCodeBase (Li et al 2023) model, which uses a decoder-only architecture with further modifications. The 3b version of these models is trained on more than 3T tokens. We chose this model as it has been developed by open source community and performs well on APR task.

granite-3b-code-base Granite models (Mishra et al. 2024) are a group of four Code LLMs released by IBM Research, ranging from 3b to 34b models. These models use the decoder blocks of the Transformer model (Vaswani et al. 2017). The pre-training process of these models consists of two phases. In the first phase, the models are trained on code data only, with the 3b model being trained on 4T of code tokens. In the second phase, the models are trained on 500B of tokens including both code and natural language tokens. In addition to casual language modeling objective, they use the Fill-In-the-Middle (FIM) (Bavarian et al. 2022) objective function during pre-training. To further align the model, they also train the models on various instruction-tuning datasets. We chose this model because it uses FIM

objective during training phase it is the most competitive on program repair benchmarka compared to other models having the same of similar size.

The instruction-tuned version of these models are publicly available. However, these versions are trained on the entire dataset that we aim to use. Hence, we will select the base checkpoints (pre-trained only) of both models and train different task-specific adapters by further instruction-tuning the base models on each task-specific dataset.

3.6 Research method for each RQ

For RQ1, we use merging with equal-weight merging (Fig. 1) and add the other four tasks one by one to be merged with APR. As the order of tasks is not important in this part, we add tasks one by one to the APR adapter and compare the results with APR adapter on the benchmark dataset (see Section 3.7). We apply this process for each of the three merging methods, weight-space averaging, TIES-Merging, and DARE-TIES, separately and report the results.

For RQ2, we use the same process as in RQ1. However, we do not add the APR adapter in the merged adapter. Instead, we test the merged adapter on APR to evaluate the generalizability of the merged adapter. We then report the results for each of the three merging methods and the combinations of the task-specific adapters.

For RQ3, we use the continual merging (Fig. 1 bottom of merging phase). For this purpose, we add the task-specific adapters one by one to build the merged adapter. In this case, the order of adding tasks is important and affects the results. We use four tasks, including APR, code development, code improvement, and test & QA as identified in Table 1. These tasks are all similar to each other in terms of being code-to-code tasks. The experiments include models when two, three, or four task-specific adapters are used in the merged adapter, and all the permutations of the tasks. We opt to remove one of the tasks to keep the number of experiments manageable according to computational resources available to us. Similar to previous RQs, we apply the same process for each of the merging techniques separately. This will show if continual merging helps improve the results compared to equal-weight averaging, and if so, it will be more effective on which merging method.

3.7 Evaluation metrics

For all experiments we report the $\text{pass}@k$ scores of the models evaluated on the HumanEvalFix benchmark (Muennighoff et al. 2023). This benchmark is created as part of the HumanEvalPack benchmarks released by the same research. They manually produced bugs in the extended version of the original HumanEval benchmark, which supports six programming languages, including Python. The HumanEvalFix benchmark contains only 164 executable code samples. Therefore, although it is categorized as an APR benchmark by its original study, it is inherently different from the source datasets used for fine-tuning the models.

Compared to previous similarity-based metrics such as BLEU (Papineni et al. 2002) and CodeBLEU (Ren et al. 2020), execution-based metrics, i.e., $\text{pass}@k$ are more reliable to capture the functional correctness of generated codes (Nijkamp et al. 2023). To calculate these scores, accessing to a dataset containing unit tests for all samples is required. In the

pass@k metric, to check the correctness of the generated code samples, one of the generated candidates should pass all the test cases of the corresponding sample in the dataset. More formally, by sampling k candidates for each record or problem of the dataset among all n generated candidates, the pass@k metric could be calculated as follows:

$$\text{pass}@k := \mathbb{E}_{\text{Problems}} \left[1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right] \quad (4)$$

where c is the number of correct samples that pass all of the unit tests of each specific problem. We report pass@1 and pass@10 similar to previous works Muennighoff et al. (2023).

Code LLMs have been shown to be sensitive to perturbations introduced in input prompts, which affects the reliability of the pass@k metric Wang et al. (2023b). Therefore, RobustPass@k (RP_s@k) (Wang et al. 2023b) was proposed to evaluate the code generation robustness of Code LLMs. Following (Wang et al. 2023b), we introduce random perturbations in the HumanEvalFix benchmark to evaluate the models in a more robust manner that reflects practical applications. To calculate the RobustPass@k metric, s random perturbations are introduced in the input prompt. n candidates are generated for each input prompt, resulting in $n \times s$ generated candidates named $f_i(x_j)$, where $1 \leq i \leq n$ and $1 \leq j \leq s$. Finally, the RobustPass@k metric is defined as below. Here, $rc_s(x) = \sum_{i=1}^n c_{i,s}(x)$ represents the worst-case correction for problem x . For each i , $c_{i,s}(x) = 1$ if and only if all $f_i(x)$ generations are correct. Following the initial study (Wang et al. 2023b), we report RP₅@1 for our experiments. However, as creating perturbations is computationally expensive and beyond our computations capacity, we only apply RobustPass@k for a few cases, and we discuss them in the Discussions in Section 5.

$$\text{RobustPass}_s@k := \mathbb{E}_x \left[1 - \frac{\binom{n - rc_s(x)}{k}}{\binom{n}{k}} \right] \quad (5)$$

We apply pairwise t-test (Kim 2015) and report the cliff's delta effect size (Cliff 1993) to examine whether the results and their effects are statistically significant. Since the scores of the generated outputs are not normally distributed, with consultation with a Statistician professor with expertise in data science and machine learning, we randomly sampled the data with a large number, i.e., 100 times, to obtain approximately normal subsets suitable for applying the t-test.

3.8 Configurations

We conduct all experiments on one NVIDIA Tesla V100 32GB GPU. The models are trained on the CommitPackFT dataset for each task and then evaluated on HumanEvalFix benchmark. We set the temperature to 0.2 when evaluating the models and generate 20 samples to calculate pass@1 and pass@10 scores. The hyperparameters are set following the previous works Muennighoff et al. (2023); Mishra et al. (2024) for training each model. Table 2

Table 2 Values of different hyperparameters used for training the models

Hyperparameter	StarCoder2	Granite
learning_rate	2e-5	1e-5
lr_scheduler_type	cosine	cosine
weight_decay	0.05	0.05
max_steps	10000	10000
warmup_steps	20	250
train_micro_batch_size	4	2
gradient_accumulation_steps	1	8
lora_rank	16	16

reports the values of these hyperparameters. For training the models using LoRA adapters, we insert LoRA layers into the decoders of the models and we apply them to `q_proj`, `o_proj`, `k_proj`, `v_proj`, `c_fc`, `c_proj` layers of the StarCoder2 and `q_proj`, `o_proj`, `k_proj`, `v_proj`, `gate_proj`, `up_proj`, `down_proj` layers of Granite model.

To generate the HumanEvalFixPerturbed dataset used in the discussions, we applied the perturbations introduced by the ReCode (Wang et al. 2023b) study for the categories of `format`, `function`, and `syntax` and used five different seeds for generating each of the datasets following their approach. We did not use the `docstring` perturbation (Wang et al. 2023b), as the `bigcode-evaluation-harness` library² we employed to evaluate the results – used in the previous studies (Muennighoff et al. 2023) – is not compatible with the prompt for the APR task. Therefore, applying perturbations on this category is not providing a ground comparison among the two metrics, `pass@k` and `RobustPass@k`.

4 Results

In this section, we report and discuss the results obtained for each RQ. Particularly, for each RQ, we report both `pass@1` and `pass@10` scores in a separate table. The tables with a light yellow background display `pass@1` scores, while those with a light blue background present the `pass@10` scores. Due to the large number of experiments for RQ3, we divide the results of the merged adapters with three and four base tasks in two categories. Additionally, for adapters having four tasks, we divide them into two tables for readability. In addition to `pass@1` and `pass@10` scores, we report the results of the statistical tests and effect size experiments conducted between each merged adapter and the APR adapter in order to investigate the importance of the results. For all experiments, we report the p-value of statistical tests with * for p-value < 0.1, ** for p-value < 0.05, and *** for p-value < 0.01. In the same vein, we report the `cliff_delta` value of the effect size experiments using S for `|cliff_delta|` >= 0.11, M for `|cliff_delta|` >= 0.28, and L for `|cliff_delta|` >= 0.43. For each model and merged adapter, the score of the best-performing merging method (row-wise) is shown in **bold**. Furthermore, the best-performing adapter for each merging method (column-wise) is underlined for both models. Note that there might be multiple adapters with the exact same scores as the highest for each merging method or model. In such cases, all are shown in **bold** or are underlined.

In all cases, we add the performance difference compared to the APR adapter in parenthesis. First, we provide the `pass@1` and `pass@10` scores of individual task-adapters on the

²<https://github.com/bigcode-project/bigcode-evaluation-harness>

APR benchmark as shown in Table 3. Interestingly, T3, which is the Misc task has highest performance on APR, even surpassing the result of APR. Test & QA on the other hand, have the lowest score among all tasks on the APR benchmark.

4.1 RQ1: Performance of merged adapters on APR

In this RQ, we merge the APR adapter with all the other task-specific adapters and evaluate the performance of the merged adapter on the APR benchmark. Tables 4 and 5 report the pass@1 and pass@10 scores of these experiments, respectively. As pass@1 scores consider the model performance with one generation of code, we mainly discuss this score for the results. However, for the completeness of the work, we also included pass@10 scores and some related discussions. Comparing the two models, StarCoder2 and Granite, the results

Table 3 Pass@1 and pass@10 scores of the individual task-specific adapters on the APR benchmark with StarCoder2 and Granite models. The highest scores are shown in **bold**

Model/Task		Tasks				
	Program Repair (T1)	Improvement (T2)	Misc (T3)	Development (T4)	Test & QA (T5)	
StarCoder2	pass@1	28.66%	27.90%	31.40%	24.66%	22.10%
	pass@10	39.11%	37.99%	41.88%	32.66%	28.90%
Granite	pass@1	16.16%	16.95%	18.63%	18.32%	13.20%
	pass@10	19.84%	21.05%	22.35%	22.79%	17.68%

Table 4 pass@1 scores of the merged adapters including T1 (APR) as one of their constituent tasks. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged does not affect the performance as they are all merged with equal weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Merged Tasks	StarCoder2			Granite		
	darc-tics	ties	weight-averaging	darc-tics	ties	weight-averaging
T1-T2	28.41% (-0.25)	29.63% (+0.97)	28.93% (+0.27)	16.43% (+0.27)	16.28% (+0.12)	16.25% (+0.09)
	S**	S**	S*	S*	S*	S*
T1-T3	30.18% (+1.52)	31.01% (+2.35)	30.15% (+1.49)	17.41% (+1.25)	17.07% (+0.91)	17.2% (+1.04)
	S**	M***	M***	M***	M***	S*
T1-T4	27.35% (-1.31)	26.89% (-1.77)	27.13% (-1.53)	17.35% (+1.19)	16.86% (+0.7)	17.35% (+1.19)
	M***	S***	M***	S**	S***	S***
T1-T5	25.24% (-3.42)	25.88% (-2.78)	25.82% (-2.84)	15.0% (-1.16)	14.73% (-1.43)	14.73% (-1.43)
	L***	L***	L***	S***	S*	L***
T1-T2-T3	30.82% (+2.16)	31.04% (+2.38)	30.95% (+2.29)	17.29% (+1.13)	17.29% (+1.13)	17.96% (+1.8)
	M***	M***	L***	S***	M***	S***
T1-T2-T4	28.2% (-0.46)	28.26% (-0.4)	28.99% (+0.33)	17.32% (+1.16)	16.98% (+0.82)	17.32% (+1.16)
	S*	S*	S**	S*	S***	S***
T1-T2-T5	27.38% (-1.28)	27.26% (-1.4)	26.49% (-2.17)	15.37% (-0.79)	15.15% (-1.01)	15.34% (-0.82)
	S***	S**	M***	S*	S*	S***
T1-T3-T4	29.27% (+0.61)	29.02% (+0.36)	28.51% (-0.15)	17.93% (+1.77)	17.8% (+1.64)	17.68% (+1.52)
	S*	S*	S*	S***	M***	M***
T1-T3-T5	28.05% (-0.61)	28.93% (+0.27)	27.99% (-0.67)	16.65% (+0.49)	16.22% (+0.06)	16.89% (+0.73)
	S***	S*	S*	S*	S*	S***
T1-T4-T5	26.16% (-2.5)	26.28% (-2.38)	25.82% (-2.84)	15.88% (-0.28)	16.01% (-0.15)	16.22% (+0.06)
	L***	M***	M***	S*	S*	S*
T1-T2-T3-T4	30.21% (+1.55)	30.06% (+1.4)	30.3% (+1.64)	17.47% (+1.31)	17.71% (+1.55)	18.14% (+1.98)
	M***	S*	M***	S**	M***	M***
T1-T2-T3-T5	29.15% (+0.49)	29.39% (+0.73)	28.05% (-0.61)	15.91% (-0.25)	16.28% (+0.12)	17.07% (+0.91)
	S*	S*	S*	S*	S*	S**
T1-T2-T4-T5	27.29% (-1.37)	27.53% (-1.13)	26.68% (-1.98)	16.49% (+0.33)	16.46% (+0.3)	17.1% (+0.94)
	S***	S*	M***	S*	S**	S*
T1-T3-T4-T5	28.69% (+0.03)	27.56% (-1.1)	27.9% (-0.76)	16.65% (+0.49)	16.74% (+0.58)	17.47% (+1.31)
	S*	S***	S*	S*	S*	S***
T1-T2-T3-T4-T5	27.8% (-0.86)	28.38% (-0.28)	27.32% (-1.34)	17.29% (+1.13)	16.65% (+0.49)	17.13% (+0.97)
	M***	S*	S*	S**	S**	S***

Table 5 pass@10 scores of the merged adapters including T1 (APR) as one of their constituent tasks. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged does not affect the performance as they are all merged with equal weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Merged Tasks	StarCoder2	ties		Granite	ties	
	dare-ties	weight-averaging	dare-ties	weight-averaging	dare-ties	weight-averaging
T1-T2	40.54% (+1.43) S***	39.88% (+0.77) S*	41.46% (+2.35) S*	19.37% (-0.47) S*	19.97% (+0.13) S*	19.51% (-0.33) S*
T1-T3	41.72% (+2.61) L***	42.24% (+3.13) L***	42.37% (+3.26) M***	21.32% (+1.48) M***	20.52% (+0.68) S***	21.29% (+1.45) M***
T1-T4	38.94% (-0.17) S*	37.78% (-1.33) M***	38.61% (-0.5) S***	21.57% (+1.73) M***	21.15% (+1.31) S***	21.65% (+1.81) M***
T1-T5	35.18% (-3.93) L***	34.43% (-4.68) L***	34.76% (-4.35) L***	19.29% (-0.55) S*	19.26% (-0.58) S*	19.49% (-0.35) S*
T1-T2-T3	41.53% (+2.42) M***	42.94% (+3.83) L***	43.12% (+4.01) M***	20.44% (+0.6) S*	21.27% (+1.43) M***	21.65% (+1.81) M***
T1-T2-T4	39.64% (+0.53) S*	40.22% (+1.11) S*	41.57% (+2.46) S***	20.44% (+0.6) S*	21.41% (+1.57) M***	21.53% (+1.69) M***
T1-T2-T5	37.91% (-1.2) S*	39.81% (+0.7) S*	37.8% (-1.31) M***	19.84% (+0.0) S*	19.48% (-0.36) S*	19.97% (+0.13) S*
T1-T3-T4	40.49% (+1.38) S**	39.39% (+0.28) S*	39.47% (+0.36) S*	21.96% (+2.12) M***	21.58% (+1.74) M***	21.16% (+1.32) S*
T1-T3-T5	38.24% (-0.87) S*	39.34% (+0.23) S*	37.25% (-1.86) L***	20.45% (+0.61) S*	19.85% (+0.01) S*	20.52% (+0.68) S**
T1-T4-T5	37.23% (-1.88) S***	38.29% (-0.82) S*	36.18% (-2.93) L***	19.76% (-0.08) S*	20.51% (+0.67) S**	20.15% (+0.31) S*
T1-T2-T3-T4	39.76% (+0.65) S***	41.14% (+2.03) S**	41.74% (+2.63) S**	20.83% (+0.99) S*	20.68% (+0.84) S***	22.45% (+2.61) L***
T1-T2-T3-T5	40.61% (+1.5) S***	40.35% (+1.24) S**	40.77% (+1.66) S*	19.18% (-0.66) S*	19.91% (+0.07) S*	20.92% (+1.08) S***
T1-T2-T4-T5	38.3% (-0.81) S*	39.09% (-0.02) S**	38.82% (-0.29) S***	21.37% (+1.53) S***	20.88% (+1.04) S***	21.11% (+1.27) S***
T1-T3-T4-T5	37.46% (-1.65) S**	38.23% (-0.88) S**	38.33% (-0.78) S***	20.55% (+0.71) S***	20.68% (+0.84) S***	21.07% (+1.23) S***
T1-T2-T3-T4-T5	39.01% (-0.1) S*	39.84% (+0.73) S*	37.31% (-1.8) M***	20.82% (+0.98) S*	19.69% (-0.15) S*	20.89% (+1.05) S***

for StarCoder2 are higher, which could reflect the ability of the StarCoder2 model for APR task.

According to the pass@1 scores in Table 4, among the 15 merged adapters with different combinations, we can observe the performance improvement over the APR adapter in less than half of the merged adapters (6 using dare-ties method, 7 using ties method, and 5 using weight-averaging method) for the first model (StarCoder2). Whereas, for the second model, this figure rises to more than half of the merged adapters (11 using dare-ties method, 12 using ties method, and 13 using weight-averaging method). In other words, a considerable number of the merged adapters evaluated with Granite model have higher scores than the APR task that is evaluated with the same model. However, for the StarCoder2 model, fewer number of the merged adapters have led to improved performance compared to the APR task. One of the reasons for this difference between different models can be rooted in the scores of other individual task-specific adapters with each model (refer to the scores reported in Table 3). Particularly, other task-specific adapters except T5, perform better than APR using the Granite model, while APR is the second best task when evaluated with the StarCoder2 model. *As a result, merging APR with other tasks is more probable to be effective when applied in Granite instead of the StarCoder2 model.* This finding is valid across all merging methods according to the pass@1 and pass@10 scores of these methods.

Comparing different merged adapters (column-wise comparison), according to the pass@1 scores of the merged adapters, T1-T2-T3 performs the best among all the other

merged adapters with each of the three merging methods when evaluated with the StarCoder2 model. The performance of the T1-T2-T3 adapter improved the APR performance, which is statistically significant and effect size is Medium to Large. However, for the Granite model, T1-T3-T4 adapter performs better than others using two merging methods, i.e., dare-ties and ties. In this case, the performance is improved compared to APR score, which is statistically significant, but the effect size is Small for dare-ties and Medium for the ties. With the weight-averaging method, T1-T2-T3-T4 performs better than other adapters, and the improvement over APR is statistically significant with Medium effect size. These results are in agreement with the results of the individual task-specific adapters, as T3 alone reaches the highest performance on the APR benchmark across both models (see Table 3). As seen above, T3 is also included in all the best-performing merged adapters.

It is also worth mentioning that while some merged adapters improve the performance of the APR, the pass@1 scores of all the merged adapters on the APR benchmark are less than the scores of the highest performing individual task-specific adapters (i.e., T3), which are 31.40% and 18.63% for StarCoder2 and Granite, respectively. That being said, it does not affect our discussions about the merging technique to improve the results, compared to the APR individual adapter. We associate the reason of high performance for the T3 adapter to the diversity of the records in its training dataset. When we investigated the results, we observed that some samples were close to program repair task. Thus, it could affect the performance of T3 on APR task in a positive way.

For the pass@10 scores, the same tasks, i.e., T1-T2-T3 and T1-T3-T4, perform better than others for StarCoder2 and Granite, respectively. Interestingly, the pass@10 scores of the T1-T2-T3 and T1-T3 merged adapters using ties and weight-averaging methods are even higher than the scores of their constituent individual task-adapters (i.e., T1, T2, and T3). Specifically, the pass@10 score of the best performing individual task-specific adapter which is 41.88% is lower than the pass@10 scores of the above-mentioned merged adapters. Despite being rare when considering the total number of the experiments conducted for different combinations of the tasks, *these results show that merging different task-specific adapters could lead to even higher performance for a specific downstream task like APR.*

We experimented with merging different *number* of task-specific adapters to construct the merge adapters as shown in Tables 4 and 5. However, we do not observe any specific pattern on the score of the merged adapters with different number of tasks involved in their merging process. In other words, increasing or decreasing the number of task-specific adapters used to build the merged adapters does not necessarily lead to performance degradation or improvement over the APR benchmark. For example, for the StarCoder2 model, the scores of the T1-T2-T3-T4-T5 adapter is higher than the scores of the T1-T5 adapter for all merging techniques, while having more constituent tasks. Conversely, the scores of the T1-T2 adapter is higher than the scores of T1-T2-T3-T4-T5 adapter in the same model for all the merging techniques. *As discussed above, we find that the type of the tasks involved in the merged adapters has a higher affect on the performance of the merged adapter rather than the number of the tasks in the merged adapter.*

When comparing the results of the merging methods in terms of the pass@1 scores (i.e., row-wise), ties method performs better than the other two methods applied on the StarCoder2 model whereas dare-ties and weight-averaging perform better if Granite is used as the base model. Surprisingly, ties performs the worst for the latter model and dare-ties and weight-averaging perform close to each other with weight-averaging being slightly better. For pass@10, a mixed results is observed for the StarCoder2 model, where weight-averaging

and ties have better scores. For the Granite model, almost the same pattern exists for the results of the $\text{pass}@10$ scores. More specifically, weight-averaging performs the best for both models with ties being close to weight-averaging for StarCoder2 and the worst for Granite.

By analyzing the merged adapters that are trained on StarCoder2 model and merged using two adapters, T1-T3 performs better than others, and after that T1-T2 is the second best-performing adapter with T1-T4 and T1-T5 coming afterwards. This suggests that T3 helps improve the performance of the APR task while the others either degrade the performance or increase it slightly (T1-T2). In the same vein, in the category of merged adapters that are created using three adapters, merging with T3 and T2 tasks lead to the highest performance improvement over other merged adapters. For the merged adapters that are built by four task-specific adapters, the combination of T2, T3, and T4 leads the highest improvement as expected. Lastly, the merged adapter of all task-specific adapters, performs as one of the middlemost adapters. These results are valid for both the $\text{pass}@1$ and $\text{pass}@10$ scores as reported in Tables 4 and 5, respectively. Again, this observation supports our discussion that the merged task plays an important role in improving the performance compared to APR adapter.

For the Granite model, almost the same pattern exists with some few changes. Specifically, since T4 performs better than T2 using this model as reported in Table 3, the merged adapter T1-T4 performs better than T1-T2. In addition, as T3 and T5 have the best and worst scores on APR, respectively, T1-T3 has the highest score and T1-T5 has the lower score among all adapters that merge two adapters. Similarly, merging T1 with T3 and T4 performs better than all merged adapters created from three tasks and merging with T2, T3, and T4 is the best-performing merged adapter among all merged adapters consisting of four tasks. Finally, the last adapters, i.e., T1-T2-T3-T4-T5 performs relatively as one of the middlemost adapters.

RQ1 Summary

Across all combinations of task-specific adapters merged using 2 to 5 constituent adapters, the scores show only marginal differences compared to the individual APR adapter. In particular, the best-performing merged adapter (T1-T2-T3) achieves a score improvement of +2.38% (with ties method) for $\text{pass}@1$ and +4.01% (with weight-averaging method) for $\text{pass}@10$ compared with the APR adapter. Similarly, the worst-performing adapter (T1-T5) shows a score degradation of -3.42% (with dare-ties method) for $\text{pass}@1$ and -4.68% (with ties method) for $\text{pass}@10$ score compared with the APR adapter. Both the highest and the lowest scores are obtained using the StarCoder2 model, having medium and large effect sizes. Among all the experiments conducted using different combinations of task-specific adapters, we can observe performance improvement on the APR benchmark in some experiments. Particularly, in 19 and 36 experiments out of 45, the $\text{pass}@1$ scores on the APR has improved for the StarCoder2 and Granite models, respectively. However, this highly depend on the type of the base model used and the constituent tasks of the merged adapter. Merging adapters using Granite as the base model is more effective in improving the performance. In addition, the individual task with higher performance than APR can lead to higher performance in the merged adapter on APR benchmark while the ones with lower performance degrade the performance of the merged adapter. Additionally, the performance of the merged adapter depends more on the task-adapter being used in the merging than the number of task-adapters being merged.

4.2 RQ2: Generalizability of Merged Adapters to APR

In this RQ, we exclude the T1 (APR) adapter when merging different combinations of task-specific adapters and evaluate the merged adapters on the APR benchmark. This will result in having merged adapters consisting of two, three, and four non-APR adapters which will ultimately lead to a set of 11 different merged adapters. The pass@1 and pass@10 scores of these merged adapters for both models via three merging methods are reported in Tables 6 and 7, respectively.

Despite of excluding the APR task when merging different task-specific adapters, we do not observe a significant performance drops on the APR benchmark when comparing with the results of RQ1. In particular, for the first model, StarCoder2, the overall performance on the APR benchmark is decreased with some merged adapters increasing the results. For the second model, Granite, the performance is increased in more cases. As mentioned earlier, the main reason of this change is related to the high performance of the other individual task-specific adapters on APR benchmark with the Granite model. Hence, the performance of the individual tasks should be considered when investigating the root causes of the performance change of the merged adapters.

According to the pass@1 scores reported in Table 6, the best performing merged adapters are T2-T3 and T3-T4 for the StarCoder2 and Granite models, respectively. For these merged adapters on both models, the results are improved compared to the APR adapter, and the difference is statistically significant with Large effect size in almost all of the merging techniques. Such results are in compliance with the scores of the individual task-specific adapters. In fact, T3 is the best performing individual task, and the performance of T2 is slightly lower than APR with the StarCoder2 model as shown in Table 3. Therefore, their merged adapter has the highest score compared to all other adapter combinations. Similarly, T3 and T4 perform better than the rest of individual adapters when used with Granite model (see Table 3). These results applies for all merging methods across both models. However, for the pass@10 scores as reported in Table 7, although the same individual tasks perform better than the rest for each model, their merged adapters do not always perform the best. Specifically, the merged adapter of T2-T3-T5 performs the best with dare-ties method for the StarCoder2 model and T2-T3-T4 performs the best with dare-ties method with Granite model. This is interesting as T5 alone has the worst performance among all task-specific adapters for StarCoder2 model and yet leads to the highest performance when merged with T2 and T3 tasks.

Considering the effectiveness of the merging methods, while ties performs slightly better than others for StarCoder2 model, it is the worst merging method when applied with Granite model. Dare-ties and weight-averaging methods perform worse than ties for the StarCoder2 model, but for the Granite model, weight-averaging is the most effective method as it performs better than others regarding both pass@1 and pass@10 scores.

Similar to the previous RQ, we do not observe much performance difference with merged adapters derived from different *number* of tasks-specific adapters. In fact, the performance of the merged adapters prone to be related with the performance of the individual tasks than the number of the tasks contributed in the merged adapters. In particular, as T3 and T2 are performing well with StarCoder2 model, their merged adapter is also the best adapter among all adapters. Similarly, for Granite, the merged adapter of T3 and T4 performs bet-

Table 6 pass@1 scores of the merged adapters consisting of non-APR task adapters. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged does not affect the performance as they are all merged with equal weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Merged Tasks	StarCoder2			Granite		
	dare-ties	ties	weight-averaging	dare-ties	ties	weight-averaging
T2-T3	30.3% (+1.64) S***	31.95% (+3.29) L***	30.67% (+2.01) L***	18.32% (+2.16) M***	17.96% (+1.8) M***	18.23% (+2.07) M***
T2-T4	27.26% (-1.4) M***	27.2% (-1.46) S***	26.62% (-2.04) S***	18.26% (+2.1) L***	18.05% (+1.89) L***	18.84% (+2.68) L***
T2-T5	25.85% (-2.81) L***	26.28% (-2.38) L***	25.82% (-2.84) M***	14.85% (-1.31) M***	14.91% (-1.25) M***	14.97% (-1.19) M***
T3-T4	27.26% (-1.4) S**	27.9% (-0.76) S***	28.99% (+0.33) S***	19.3% (+3.14) L***	18.84% (+2.68) L***	19.09% (+2.93) L***
T3-T5	26.01% (-2.65) M***	26.83% (-1.83) M***	26.37% (-2.29) M***	16.77% (+0.61) S**	16.55% (+0.39) S*	17.41% (+1.25) S***
T4-T5	24.42% (-4.24) L***	24.79% (-3.87) L***	24.91% (-3.75) L***	16.34% (+0.18) S*	15.82% (-0.34) S*	16.43% (+0.27) S*
T2-T3-T4	29.27% (+0.61) S*	29.12% (+0.46) S*	29.09% (+0.43) S**	18.48% (+2.32) L***	18.78% (+2.62) L***	18.99% (+2.83) L***
T2-T3-T5	29.02% (+0.36) S*	28.48% (-0.18) S*	27.77% (-0.89) S*	17.04% (+0.88) S**	16.55% (+0.39) S**	16.8% (+0.64) S*
T2-T4-T5	25.09% (-3.57) L***	25.88% (-2.78) L***	26.16% (-2.5) M***	16.74% (+0.58) S*	16.25% (+0.09) S*	17.5% (+1.34) M***
T3-T4-T5	27.8% (-0.86) S*	27.1% (-1.56) M***	26.37% (-2.29) M***	17.23% (+1.07) S***	17.26% (+1.1) M***	17.8% (+1.64) M***
T2-T3-T4-T5	27.5% (-1.16) S***	27.84% (-0.82) S*	27.29% (-1.37) S***	16.92% (+0.76) S**	17.47% (+1.31) M***	17.77% (+1.61) M***

Table 7 pass@10 scores of the merged adapters consisting of non-APR task adapters. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged does not affect the performance as they are all merged with equal weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Merged Tasks	StarCoder2			Granite		
	dare-ties	ties	weight-averaging	dare-ties	ties	weight-averaging
T2-T3	39.48% (+0.37) S*	41.79% (+2.68) M***	40.79% (+1.68) M***	21.53% (+1.69) S**	20.86% (+1.02) S**	21.78% (+1.94) M***
T2-T4	37.09% (-2.02) L***	37.89% (-1.22) S**	35.94% (-3.17) M***	22.27% (+2.43) M***	21.98% (+2.14) M***	22.63% (+2.79) L***
T2-T5	35.62% (-3.49) L***	36.49% (-2.62) M***	35.17% (-3.94) L***	19.58% (-0.26) S*	18.78% (-1.06) M***	18.83% (-1.01) S*
T3-T4	37.5% (-1.61) M***	39.21% (+0.1) S**	40.47% (+1.36) S***	21.88% (+2.04) S***	21.9% (+2.06) M***	22.1% (+2.26) L***
T3-T5	35.46% (-3.65) L***	37.0% (-2.11) S***	35.78% (-3.33) M***	21.4% (+1.56) S***	20.49% (+0.65) S*	21.58% (+1.74) M***
T4-T5	32.54% (-6.57) L***	34.33% (-4.78) L***	33.25% (-5.86) L***	20.82% (+0.98) S*	19.5% (-0.34) S*	20.58% (+0.74) S***
T2-T3-T4	39.45% (+0.34) S*	41.02% (+1.91) M***	40.71% (+1.6) M***	22.41% (+2.57) L***	21.59% (+1.75) M***	22.2% (+2.36) L***
T2-T3-T5	39.6% (+0.49) S*	39.17% (+0.06) S*	39.28% (+0.17) S**	20.52% (+0.68) S*	19.79% (-0.05) S*	20.08% (+0.24) S*
T2-T4-T5	33.4% (-5.71) L***	35.64% (-3.47) M***	35.96% (-3.15) M***	20.83% (+0.99) S*	20.33% (+0.49) S*	21.41% (+1.57) M***
T3-T4-T5	37.29% (-1.82) M***	37.58% (-1.53) S**	36.01% (-3.1) M***	20.89% (+1.05) S**	20.76% (+0.92) S**	21.42% (+1.58) M***
T2-T3-T4-T5	38.27% (-0.84) M***	38.2% (-0.91) S*	36.93% (-2.18) S*	20.5% (+0.66) S*	20.56% (+0.72) S*	21.63% (+1.79) M***

ter than other adapters as expected. Note that in this RQ, merging four adapters reduces the performance compared to most of the merged adapter with two or three adapters.

RQ2 Summary

The score differences compared to the APR adapter are significant with large effect size; the best-performing adapter (T2-T3) achieves a +3.29% (with ties method) improvement in pass@1 with the StarCoder2 model. For pass@10, the best-performing adapter shows a +2.79% improvement over the APR adapter (T2-T4 using Granite model and with weight-averaging method). In contrast, the worst-performing adapter (T4-T5) shows a score degradation of -4.24% for pass@1 and -6.57% for pass@10, both using the StarCoder2 model and dare-ties method. This result suggests that merging adapters with consideration of the Tasks in the merged model can significantly affect the performance.

Similar to RQ1, the performance of the merged adapters, whether it rises or falls, depends on the performance of the individual task-specific adapters on the APR benchmark. The individual adapters' performance has the greatest influence on the final score of the merged adapter. In other words, if the individual task adapters perform well, the merged adapter is unlikely to experience a significant performance drop. Consequently, we observe that merged adapters originating from non-APR tasks can still perform well on the APR benchmark, provided their constituent adapters achieve strong APR scores. In some cases, the merged adapter even outperforms the APR-specific adapter, demonstrating the generalizability of the merging approach. This performance improvement over APR-specific adapter is observed in 8 and 29 pass@1 scores out of 33 experiments for StarCoder2 and Granite models, respectively.

4.3 RQ3: Performance of continual merging

In this RQ, we are specifically interested in investigating the effect of the order of the adapters added in the merged adapter on the performance of the merged adapter on the APR benchmark. As aforementioned, the sequential addition of task-specific adapters will lead to higher weight of the adapters which were merged lastly. As such, we can argue that the performance of the merged adapter is close to that of the last adapters in the continual merging process. Please note that due to resource restrictions, we only experiment with 4 tasks (i.e., APR, Development, Improvement, and Test & QA) to be merged with varying orders. We opted for excluding the Misc adapter as after investigating the dataset, we found that this category did not represent one single type of code change and could not help much in terms of the interpretation of the results. However, we provide the related experiments including Misc (T3) in the Appendix. Note that the permutation of two tasks would remain the same under continual merging, and are already presented in RQ1. The pass@1 results of this RQ are shown in Tables 8 and 9. Tables 10 and 11 represent the results for pass@10. Note that for pass@1, most results are statistically significant, with mainly effect sizes of Medium or Large for StarCoder2 and varying effect sizes for Granite.

Among these four tasks, excluding T3, APR (T1) is the best performing task with StarCoder2 and Development (T4) with the Granite model (see Table 3). Based on this, we

Table 8 pass@1 scores of the merged adapters using continual merging approach built from three task-specific adapters. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-av eraging
T1, T2, T4	T1-T2-T4	27.32% (-1.34)	27.35% (-1.31)	27.5% (-1.16)	17.84% (+1.68)	17.47% (+1.31)	17.99% (+1.83)
		S*	S**	S***	M***	M***	S***
	T1-T4-T2	28.29% (-0.37)	28.81% (+0.15)	28.78% (+0.12)	17.47% (+1.31)	17.32% (+1.16)	17.23% (+1.07)
		S*	S*	S*	M***	S***	S***
	T2-T1-T4	28.17% (-0.49)	27.35% (-1.31)	27.5% (-1.16)	17.71% (+1.55)	17.47% (+1.31)	17.99% (+1.83)
		S*	S*	S***	M***	S**	M***
	T2-T4-T1	29.21% (+0.55)	28.96% (+0.3)	28.78% (+0.12)	17.07% (+0.91)	16.55% (+0.39)	16.86% (+0.7)
		S*	S*	S*	S*	S*	S*
	T4-T1-T2	28.11% (-0.55)	28.81% (+0.15)	28.78% (+0.12)	17.47% (+1.31)	17.32% (+1.16)	17.23% (+1.07)
		S**	S*	S*	M***	S***	S***
	T4-T2-T1	29.3% (+0.64)	28.96% (+0.3)	28.78% (+0.12)	17.01% (+0.85)	16.55% (+0.39)	16.86% (+0.7)
		S**	S**	S*	S***	S*	S**
T1, T2, T5	T1-T2-T5	25.73% (-2.93)	25.88% (-2.78)	25.37% (-3.29)	15.21% (-0.95)	14.24% (-1.92)	15.24% (-0.92)
		L***	M***	L***	S***	M***	S***
	T1-T5-T2	27.9% (-0.76)	28.11% (-0.55)	27.59% (-1.07)	16.28% (+0.12)	15.61% (-0.55)	15.91% (-0.25)
		S*	S*	S***	S*	S*	S*
	T2-T1-T5	25.52% (-3.14)	25.88% (-2.78)	25.37% (-3.29)	14.36% (-1.8)	14.24% (-1.92)	15.24% (-0.92)
		L***	M***	L***	S***	M***	S**
	T2-T5-T1	27.68% (-0.98)	28.2% (-0.46)	27.41% (-1.25)	15.7% (-0.46)	15.58% (-0.58)	15.58% (-0.58)
		S***	S*	M***	S*	S*	S*
	T5-T1-T2	27.84% (-0.82)	28.11% (-0.55)	27.59% (-1.07)	15.82% (-0.34)	15.61% (-0.55)	15.91% (-0.25)
		S*	S*	S***	S*	S*	S*
	T5-T2-T1	28.11% (-0.55)	28.2% (-0.46)	27.41% (-1.25)	15.73% (-0.43)	15.58% (-0.58)	15.58% (-0.58)
		S*	S*	S***	S*	S*	S*
T1, T4, T5	T1-T4-T5	25.43% (-3.23)	24.63% (-4.03)	25.43% (-3.23)	15.49% (-0.67)	15.34% (-0.82)	15.55% (-0.61)
		L***	L***	L***	S*	S*	S**
	T1-T5-T4	26.07% (-2.59)	25.09% (-3.57)	25.95% (-2.71)	16.92% (+0.76)	16.55% (+0.39)	17.26% (+1.1)
		L***	L***	L***	S***	S*	S***
	T4-T1-T5	24.73% (-3.93)	24.63% (-4.03)	25.43% (-3.23)	15.7% (-0.46)	15.34% (-0.82)	15.55% (-0.61)
		L***	L***	L***	S*	S**	S*
	T4-T5-T1	27.07% (-1.59)	25.73% (-2.93)	26.46% (-2.2)	16.49% (+0.33)	16.07% (-0.09)	16.1% (-0.06)
		S***	M***	M***	S**	S*	S*
	T5-T1-T4	26.28% (-2.38)	25.09% (-3.57)	25.95% (-2.71)	16.55% (+0.39)	16.55% (+0.39)	17.26% (+1.1)
		L***	L***	L***	S**	S*	S*
	T5-T4-T1	26.8% (-1.86)	25.73% (-2.93)	26.46% (-2.2)	16.37% (+0.21)	16.07% (-0.09)	16.1% (-0.06)
		S***	L***	M***	S*	S*	S*

observe that the scores of the sequentially merged adapters are aligned with the performance of the individual task-specific adapters. Looking at the pass@1 scores of merging 3 tasks as reported in Table 8, we can verify that the last adapters (the rightmost ones) have a significant effect on the performance of the merged adapter. For example, for the base tasks of T1, T2, T4, the merged adapters with T2-T4-T1 and T4-T2-T1 orders perform better than the rest via all merging methods for the StarCoder2 model. Likewise, for Granite, the scores of the merged tasks with T1-T2-T4 and T2-T1-T4 orders are higher than those of the others with the same set of tasks (T1, T2, and T4).

In this vein, the performance of the merged tasks consisting of three task-specific adapters built from other base tasks, follows a similar pattern. Particularly, for the base tasks of T1, T2, and T5, the merged adapters with either T1 or T2 as the last added task perform better than those two with T5 as the last task (for the StarCoder2 model). This happens as the score of the individual task-specific adapter of T5 is the lowest on the APR benchmark via StarCoder2 model. Likewise, for the merged adapters of T1, T4, and T5 the same results are obtained when T5 is added in the last step. These trends are valid for the pass@10 scores of the merged adapters as reported in Table 10.

Table 9 pass@1 scores of the merged adapters using continual merging approach built from four task-specific adapters. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-averaging
T1-T2-T4-T5	T1-T2-T4-T5	25.03% (-3.63)	24.33% (-4.33)	25.73% (-2.93)	15.85% (-0.31)	15.4% (-0.76)	16.28% (+0.12)
		L***	L***	L***	S*	S*	S*
T1-T2-T5-T4	T1-T2-T5-T4	25.21% (-3.45)	25.55% (-3.11)	25.67% (-2.99)	17.04% (+0.88)	16.43% (+0.27)	17.1% (+0.94)
		L***	M***	L***	S***	S*	S*
T1-T4-T2-T5	T1-T4-T2-T5	25.43% (-3.23)	25.49% (-3.17)	25.73% (-2.93)	14.66% (-1.5)	14.82% (-1.34)	15.82% (-0.34)
		L***	L***	L***	S**	M***	S*
T1-T4-T5-T2	T1-T4-T5-T2	26.31% (-2.35)	27.29% (-1.37)	27.16% (-1.5)	17.13% (+0.97)	16.13% (-0.03)	16.37% (+0.21)
		M***	S*	M***	S**	S*	S*
T1-T5-T2-T4	T1-T5-T2-T4	25.85% (-2.81)	26.46% (-2.2)	26.43% (-2.23)	17.65% (+1.49)	16.95% (+0.79)	17.87% (+1.71)
		L***	M***	L***	L***	S*	M***
T1-T5-T4-T2	T1-T5-T4-T2	27.47% (-1.19)	27.96% (-0.7)	27.38% (-1.28)	17.56% (+1.4)	16.77% (+0.61)	17.01% (+0.85)
		S*	S*	M***	M***	S***	S***
T2-T1-T4-T5	T2-T1-T4-T5	25.46% (-3.2)	24.33% (-4.33)	25.73% (-2.93)	15.85% (-0.31)	15.4% (-0.76)	16.28% (+0.12)
		L***	L***	L***	S*	S***	S*
T2-T1-T5-T4	T2-T1-T5-T4	25.76% (-2.9)	25.55% (-3.11)	25.67% (-2.99)	17.23% (+1.07)	16.43% (+0.27)	17.1% (+0.94)
		L***	M***	L***	S***	S*	M***
T2-T4-T1-T5	T2-T4-T1-T5	25.4% (-3.26)	25.3% (-3.36)	25.55% (-3.11)	14.7% (-1.46)	14.7% (-1.46)	15.43% (-0.73)
		L***	L***	L***	S***	M***	S*
T2-T4-T5-T1	T2-T4-T5-T1	27.04% (-1.62)	27.35% (-1.31)	27.13% (-1.53)	15.79% (-0.37)	15.46% (-0.7)	16.13% (-0.03)
		M***	S**	M***	S*	S*	S*
T2-T5-T1-T4	T2-T5-T1-T4	26.77% (-1.89)	27.16% (-1.5)	26.52% (-2.14)	17.44% (+1.28)	17.26% (+1.1)	17.16% (+1.0)
		M***	S*	M***	M***	S***	S***
T2-T5-T4-T1	T2-T5-T4-T1	27.59% (-1.07)	27.2% (-1.46)	27.32% (-1.34)	16.62% (+0.46)	16.43% (+0.27)	16.8% (+0.64)
		S*	S*	S***	S**	S*	S*
T4-T1-T2-T5	T4-T1-T2-T5	25.73% (-2.93)	25.49% (-3.17)	25.73% (-2.93)	14.94% (-1.22)	14.82% (-1.34)	15.82% (-0.34)
		L***	M***	L***	S**	M***	S*
T4-T1-T5-T2	T4-T1-T5-T2	27.01% (-1.65)	27.29% (-1.37)	27.16% (-1.5)	16.46% (+0.3)	16.13% (-0.03)	16.37% (+0.21)
		M***	S***	S***	S**	S*	S*
T4-T2-T1-T5	T4-T2-T1-T5	25.4% (-3.26)	25.3% (-3.36)	25.55% (-3.11)	15.37% (-0.79)	14.7% (-1.46)	15.43% (-0.73)
		L***	L***	L***	S*	M***	S*
T4-T2-T5-T1	T4-T2-T5-T1	28.02% (-0.64)	27.35% (-1.31)	27.13% (-1.53)	15.98% (-0.18)	15.46% (-0.7)	16.13% (-0.03)
		S*	S*	M***	S*	S**	S*
T4-T5-T1-T2	T4-T5-T1-T2	28.69% (+0.03)	28.38% (-0.28)	28.26% (-0.4)	16.89% (+0.73)	16.34% (+0.18)	16.86% (+0.7)
		S*	S*	S*	S***	S*	S**
T4-T5-T2-T1	T4-T5-T2-T1	28.57% (-0.09)	28.69% (+0.03)	29.02% (+0.36)	15.95% (-0.21)	16.13% (-0.03)	16.62% (+0.46)
		S*	S*	S*	S*	S*	S*
T5-T1-T2-T4	T5-T1-T2-T4	27.26% (-1.4)	26.46% (-2.2)	26.43% (-2.23)	17.87% (+1.71)	16.95% (+0.79)	17.87% (+1.71)
		S***	M***	L***	L***	S*	M***
T5-T1-T4-T2	T5-T1-T4-T2	28.08% (-0.58)	27.96% (-0.7)	27.38% (-1.28)	17.5% (+1.34)	16.77% (+0.61)	17.01% (+0.85)
		S*	S*	S*	M***	S**	S*
T5-T2-T1-T4	T5-T2-T1-T4	27.41% (-1.25)	27.16% (-1.5)	26.52% (-2.14)	17.68% (+1.52)	17.26% (+1.1)	17.16% (+1.0)
		S***	S**	L***	M***	S***	S***
T5-T2-T4-T1	T5-T2-T4-T1	27.62% (-1.04)	27.2% (-1.46)	27.32% (-1.34)	16.52% (+0.36)	16.43% (+0.27)	16.8% (+0.64)
		S**	S**	S**	S*	S*	S*
T5-T4-T1-T2	T5-T4-T1-T2	28.32% (-0.34)	28.38% (-0.28)	28.26% (-0.4)	16.74% (+0.58)	16.34% (+0.18)	16.86% (+0.7)
		S*	S*	S*	S***	S*	S**
T5-T4-T2-T1	T5-T4-T2-T1	28.41% (-0.25)	28.69% (+0.03)	29.02% (+0.36)	16.01% (-0.15)	16.13% (-0.03)	16.62% (+0.46)
		S*	S**	S*	S*	S*	S**

When merging four task-specific adapters, we can still observe that the low-performing tasks such as T5 degrade the performance of the merged adapters if they are merged in the last steps of the continual merging. Particularly, almost all of the merged adapters having T5 as their last or second last adapter, have the pass@1 score of close to 25% via all merging methods with StarCoder2 model. Similar poor results are obtained for the Granite model as reported in Table 9. On the other hand, the merged adapters having T1 and T2 as the rightmost adapters, perform better than the others in most cases. For example, for StarCoder2 model, T4-T5-T1-T2, T4-T5-T2-T1, and T5-T4-T2-T1 are the best performing tasks. For the Granite model, the best performing tasks are T5-T1-T2-T4, T5-T2-T1-T4, T2-T5-T1-

Table 10 pass@10 scores of the merged adapters using continual merging approach built from three task-specific adapters. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA).

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-averaging
T1,T2,T4	T1-T2-T4	40.25% (+1.14) M***	39.12% (+0.01) S*	38.5% (-0.61) S*	22.31% (+2.47) L***	21.38% (+1.54) S*	21.57% (+1.73) M***
	T1-T4-T2	39.32% (+0.21) S**	39.57% (+0.46) S*	39.21% (+0.1) S***	21.84% (+2.0) M***	22.09% (+2.25) M***	21.73% (+1.89) M***
	T2-T1-T4	39.15% (+0.04) S*	39.12% (+0.01) S*	38.5% (-0.61) S*	21.56% (+1.72) M***	21.38% (+1.54) S**	21.57% (+1.73) M***
	T2-T4-T1	39.78% (+0.67) S***	39.98% (+0.87) S*	40.75% (+1.64) M***	22.38% (+2.54) S*	20.6% (+0.76) S*	21.25% (+1.41) M***
	T4-T1-T2	39.81% (+0.7) S*	39.57% (+0.46) S*	39.21% (+0.1) S*	21.41% (+1.57) M***	22.09% (+2.25) S***	21.73% (+1.89) M***
	T4-T2-T1	41.51% (+2.4) L***	39.98% (+0.87) S*	40.75% (+1.64) M***	22.11% (+2.27) M***	20.6% (+0.76) S*	21.25% (+1.41) S***
	T1-T2-T5	34.9% (-4.21) L***	34.99% (-4.12) L***	35.05% (-4.06) L***	19.13% (-0.71) S*	18.99% (-0.85) S***	19.44% (-0.4) S*
	T1-T5-T2	39.26% (+0.15) S*	39.49% (+0.38) S*	39.74% (+0.63) S***	20.43% (+0.59) S*	19.47% (-0.37) S***	20.19% (+0.35) S***
	T2-T1-T5	34.71% (-4.4) L***	34.99% (-4.12) L***	35.05% (-4.06) L***	19.45% (-0.39) S*	18.99% (-0.85) S***	19.44% (-0.4) S*
	T2-T5-T1	40.32% (+1.21) S***	41.06% (+1.95) M***	39.89% (+0.78) S***	19.91% (+0.07) S*	19.84% (+0.0) S**	20.3% (+0.46) S*
T1,T2,T5	T5-T1-T2	38.82% (-0.29) S*	39.49% (+0.38) S*	39.74% (+0.63) S***	20.68% (+0.84) S*	19.47% (-0.37) S***	20.19% (+0.35) S**
	T5-T2-T1	39.88% (+0.77) S***	41.06% (+1.95) M***	39.89% (+0.78) S***	20.02% (+0.18) S*	19.84% (+0.0) S*	20.3% (+0.46) S**
	T1-T4-T5	35.16% (-3.95) L***	33.71% (-5.4) L***	35.12% (-3.99) L***	20.22% (+0.38) S*	19.78% (-0.06) S*	19.84% (+0.0) S*
	T1-T5-T4	36.34% (-2.77) S***	34.11% (-5.0) L***	37.06% (-2.05) S***	21.25% (+1.41) M***	20.68% (+0.84) S*	21.31% (+1.47) M***
	T4-T1-T5	33.7% (-5.41) L***	33.71% (-5.4) L***	35.12% (-3.99) M***	20.38% (+0.54) S*	19.78% (-0.06) S*	19.84% (+0.0) S*
	T4-T5-T1	37.88% (-1.23) S*	37.16% (-1.95) S***	37.19% (-1.92) S*	21.03% (+1.19) S**	20.07% (+0.23) S*	19.77% (-0.07) S*
	T5-T1-T4	35.19% (-3.92) M***	34.11% (-5.0) L***	37.06% (-2.05) S**	20.81% (+0.97) S*	20.68% (+0.84) S*	21.31% (+1.47) M***
	T5-T4-T1	38.03% (-1.08) S*	37.16% (-1.95) M***	37.19% (-1.92) S**	19.83% (-0.01) S*	20.07% (+0.23) S*	19.77% (-0.07) S*

T4, and T1-T5-T2-T4 in all of which the T4 task exists as the last one. Recall that for this model, T4 has the best pass@1 score compared to the rest of the individual task-specific adapters excluding T3. Almost the same tasks performed the best based on the pass@10 scores across both models and all three merging methods as observed from Table 11.

Interestingly, we still observe that in several cases the scores of the merged adapters are higher than the individual scores of their base tasks. In fact, for the StarCoder2 model, all the merged adapters that have positive performance difference as reported in parenthesis in front of their real scores, have higher performance than that of their individual task adapters. Please note that this performance difference reported in all tables is against T1 (APR) task; hence, as T1 has the highest performance among all the individual task adapters (excluding T3) for the StarCoder2 model, the previous interpretation is derived. This pattern happens for both pass@1 and pass@10 scores of the first model. For the second model, however, as T4 is the best performing task instead of T1, there exist more considerable number of merged adapters that have higher performance compared to T1 task for both pass@1 and pass@10 scores.

Table 11 pass@10 scores of the merged adapters using continual merging approach built from four task-specific adapters. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-av eraging
T1,T2,T4,T5	T1-T2-T4-T5	34.03% (-5.08)	34.07% (-5.04)	35.46% (-3.65)	20.51% (+0.67)	19.92% (+0.08)	20.56% (+0.72)
		L***	L***	L***	S**	S*	S**
		35.09% (-4.02)	35.79% (-3.32)	36.31% (-2.8)	20.96% (+1.12)	20.36% (+0.52)	21.11% (+1.27)
T1,T2,T4,T5	T1-T2-T5-T4	L***	L***	S***	S***	S*	S***
		34.17% (-4.94)	34.92% (-4.19)	35.77% (-3.34)	19.58% (-0.26)	19.14% (-0.7)	20.51% (+0.67)
		L***	L***	L***	S*	S*	S**
T1,T2,T4,T5	T1-T4-T5-T2	35.96% (-3.15)	37.97% (-1.14)	38.15% (-0.96)	20.7% (+0.86)	20.98% (+1.14)	20.28% (+0.44)
		L***	S***	S*	S***	S*	S*
		36.99% (-2.12)	37.62% (-1.49)	35.7% (-3.41)	21.59% (+1.75)	20.98% (+1.14)	21.99% (+2.15)
T1,T2,T4,T5	T1-T5-T2-T4	S**	S**	M***	M***	S*	L***
		38.48% (-0.63)	37.87% (-1.24)	37.68% (-1.43)	21.46% (+1.62)	20.69% (+0.85)	20.86% (+1.02)
		S*	S***	S*	S***	S*	S**
T1,T2,T4,T5	T2-T1-T4-T5	34.34% (-4.77)	34.07% (-5.04)	35.46% (-3.65)	20.53% (+0.69)	19.92% (+0.08)	20.56% (+0.72)
		L***	L***	L***	S*	S*	S**
		35.99% (-3.12)	35.79% (-3.32)	36.31% (-2.8)	20.86% (+1.02)	20.36% (+0.52)	21.11% (+1.27)
T1,T2,T4,T5	T2-T1-T5-T4	M***	L***	S***	S**	S*	S***
		35.54% (-3.57)	34.9% (-4.21)	35.22% (-3.89)	20.02% (+0.18)	19.0% (-0.84)	19.85% (+0.01)
		L***	L***	L***	S*	M***	S*
T1,T2,T4,T5	T2-T4-T5-T1	37.78% (-1.33)	38.94% (-0.17)	37.26% (-1.85)	19.97% (+0.13)	19.75% (-0.09)	20.76% (+0.92)
		S*	S*	S**	S*	S*	S**
		38.17% (-0.94)	38.9% (-0.21)	36.61% (-2.5)	21.45% (+1.61)	21.29% (+1.45)	21.09% (+1.25)
T1,T2,T4,T5	T2-T5-T1-T4	S*	S*	S***	M***	S*	S***
		39.83% (+0.72)	40.27% (+1.16)	38.46% (-0.65)	20.76% (+0.92)	20.57% (+0.73)	21.23% (+1.39)
		S***	S*	S*	S*	S*	S***
T1,T2,T4,T5	T4-T1-T2-T5	33.94% (-5.17)	34.92% (-4.19)	35.77% (-3.34)	19.02% (-0.82)	19.14% (-0.7)	20.51% (+0.67)
		L***	L***	M***	S*	S***	S*
		38.76% (-0.35)	37.97% (-1.14)	38.15% (-0.96)	20.56% (+0.72)	20.98% (+1.14)	20.28% (+0.44)
T1,T2,T4,T5	T4-T1-T5-T2	S*	S*	S*	S*	S*	S*
		34.12% (-4.99)	34.9% (-4.21)	35.22% (-3.89)	20.44% (+0.6)	19.0% (-0.84)	19.85% (+0.01)
		L***	L***	L***	M***	S***	S*
T1,T2,T4,T5	T4-T2-T5-T1	40.1% (+0.99)	38.94% (-0.17)	37.26% (-1.85)	20.7% (+0.86)	19.75% (-0.09)	20.76% (+0.92)
		S***	S*	S*	S**	S**	S***
		37.24% (-1.87)	37.62% (-1.49)	35.7% (-3.41)	21.52% (+1.68)	20.98% (+1.14)	21.99% (+2.15)
T1,T2,T4,T5	T5-T1-T2-T4	S*	S***	L***	M***	S*	M***
		39.48% (+0.37)	37.87% (-1.24)	37.68% (-1.43)	21.44% (+1.6)	20.69% (+0.85)	20.86% (+1.02)
		S*	S***	S*	S**	S*	M***
T1,T2,T4,T5	T5-T2-T1-T4	38.68% (-0.43)	38.9% (-0.21)	36.61% (-2.5)	21.43% (+1.59)	21.29% (+1.45)	21.09% (+1.25)
		S*	S*	M***	M***	S**	S***
		40.1% (+0.99)	40.27% (+1.16)	38.46% (-0.65)	20.7% (+0.86)	20.57% (+0.73)	21.23% (+1.39)
T1,T2,T4,T5	T5-T2-T4-T1	M***	S**	S*	S***	S*	M***
		40.39% (+1.28)	39.16% (+0.05)	40.23% (+1.12)	20.56% (+0.72)	19.98% (+0.14)	21.15% (+1.31)
		S***	S*	M***	S*	S*	M***
T1,T2,T4,T5	T5-T4-T2-T1	41.05% (+1.94)	42.12% (+3.01)	41.87% (+2.76)	20.78% (+0.94)	20.13% (+0.29)	20.52% (+0.68)
		M***	L***	L***	S*	S*	S*

Comparing the results of this RQ with the results of the first RQ, we can observe that continual merging can lead to better performance on APR benchmark if the order of the tasks are selected carefully. Figure 2 illustrates the comparison between pass@1 scores of the T1, T2, and T4 tasks merged using equal-weight merging (RQ1) and continual merging along with the APR adapter, using StarCoder2 model. In all plots, the APR plot is shown in blue and the red plot refers to the equal-weight merging adapter (RQ1). As it is observed, when using dare-ties or ties merging methods, we can achieve better performance when the order of the merged adapters is optimal. The performance of the equal-weight merged adapter of these adapters is slightly higher than the performance of all the merged adapters using continual merging.




Fig. 2 Bar plot of pass@1 scores obtained by merged adapters of T1, T2, and T4 using continual merging (RQ3) and equal-weight merging (RQ1) approaches using StarCoder2 model. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Figure 3 shows pass@1 score for the adapters derived by T1, T4, and T5 tasks on top of StarCoder2 model. Unlike the previous tasks, here we can observe that the equal-weight merging approach only performs better than continual merging with the ties method. With the two other merging methods, we can again achieve better performance in continual merging using two orders that contain the APR (T1) task in the last position.

Similarly, with the Granite model, we achieve better performance using continual merging compared to equal-weight merging when an efficient order in continual merging is used. Figures 4 and 5 represent the associated plots. Using Granite model, the performance is improved with continual merging with any of the merging methods.

As the order of the base tasks plays a significant role in the results of the merged adapter, knowing with which order we can obtain a superior result can help reduce the trials to find the best order of merged adapters. According to the results reported in Table 8, having a set of tasks, a greedy approach to find the best order of tasks that can perform better than others is to sort the best performing tasks based on their individual scores (reported in Table 3) and merge them in ascending order; meaning that the tasks with higher scores should be merged in the last steps. For example, for the tasks of T1, T2, and T4, their ascending order based on their individual performances is T4-T2-T1 for the StarCoder2 model. Consequently, the merged adapter of T4-T2-T1 performs the best




Fig. 5 Bar plot of pass@1 scores obtained by merged adapters of T1, T4, and T5 using continual merging (RQ3) and equal-weight merging (RQ1) approaches using Granite model. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

among all the merged adapters of the same base tasks. Similarly, for the tasks of T1, T2, and T5, the best order is T5-T2-T1 which has the highest performance over the merged adapters of the same tasks. Also, for T1, T4, and T5 the best order is T5-T4-T1, which is again the best performing merged adapter in its category. Note that these results are consistent across all merging methods. However, when merging four tasks together, this greedy approach does not always lead to the best performance of the adapters. As it is observed in Table 9, the best performing adapters are T4-T5-T1-T2 and T4-T5-T2-T1 and none of them has the best order based on the individual results.

For the second model, i.e., granite-2b-code-base, the best order of tasks is T5-T1-T2-T4, but this order does not always perform the best when merging all these adapters together except for the both pass@1 and pass@10 scores of the weight-averaging method. Unlike the previous model, for the merged adapters with three base adapters the greedy approach does not always produce merged adapters with the highest scores. For example, when merging T1, T2, and T5 adapters, the best performing adapter is T1-T5-T2 which does not have the best order based on their individual performances.

For pass@1 and pass@10, similar observations are noted when including Misc (T3). See results in the Appendix.

RQ3 Summary

In this RQ, the performance differences between merged adapters and the APR adapter are still observed. In particular, the best-performing adapters are built from three tasks (T1, T2, and T4), with the orders T1-T2-T4 and T2-T1-T4 achieving a pass@1 improvement of +1.83% (with weight-averaging method) over the APR adapter using the Granite model. For pass@10 , the best-performing adapter is T4-T5-T2-T1, with a +3.01% (with ties method) improvement over the APR adapter, obtained using the StarCoder2 model. The worst-performing adapters are T1-T2-T4-T5 and T2-T1-T4-T5, which show a pass@1 degradation of -4.33% (with ties method) compared to the APR adapter using the StarCoder2 model. For pass@10 , the worst-performing adapter is T4-T1-T5, with a -5.41% (with dare-ties method) degradation compared to the APR adapter. While the improvements for best cases have small or medium effect size, the worst performing ones differ largely with the performance of the APR task.

In the continual merging approach, the order of merging task-specific adapters affects the performance of the merged adapters significantly. It is consistently observed that merging the best performing individual adapters in the last iterations lead to higher overall performance of the merged adapter on the APR benchmark. This is related with the weight influence of adapters as the lastly added adapters have higher weights in the merged adapter. Such observation happens among all merging methods and both models. On the other hand, merging the low-performance adapters in the last iterations can degrade the performance of the merged adapter.

Additionally, for most of the merged task-specific adapters, we observe an improved performance when using continual merging, considering the best order of the adapter, compared to their equal-weight averaging used in RQ1. This shows that having different weights for the task-specific adapters and considering the order of the adapters in the merged adapter can affect the performance of the merged adapter.

5 Discussion

This section provides a more detailed analysis of the experiments conducted on merging task-specific adapters. Specifically, we present examples from the HumanEvalFix benchmark, highlighting cases where individual adapters either succeeded or failed to solve the problems. Additionally, we compare the efficiency of the continual merging approach with the equal-weight merging approach, and evaluate the robustness of both individual and merged adapters under input perturbations. Finally, we analyze the Fraction of Sign Difference in the weight parameters of the adapters.

5.1 Example Cases

As mentioned earlier, in some cases, we observed that merged adapters could outperform the individual task-specific adapters on the APR benchmark. For example, the pass@10 score of the T1-T2-T3 merged adapter is higher than that of any individual task-specific adapter. While there were many samples in the evaluation dataset that all these adapters

could solve, there were a few examples that only the merged version was able to solve. As shown in Table 12, the problem is solved exclusively by the merged adapter and remains unsolved when using the individual task adapters. This suggests that merging different task adapters can be beneficial for solving certain samples.

5.2 Robustness evaluation of task-specific adapters

Code LLMs have been shown to be sensitive to small perturbations in input prompts. In real-world applications, such issues with prompts are common. However, most current benchmarks used to evaluate Code LLMs do not adequately capture the robustness capabilities of these models. A recent study (Wang et al. 2023b) proposed a framework for introducing intentional perturbations to the input prompts, thereby enabling the evaluation of model robustness, primarily for the code generation task.

Following this work, we extend their framework to the automated program repair task by applying their perturbations to the HumanEvalFix dataset, resulting in a new dataset called **HumanEvalFixPerturbed**. This dataset is used to evaluate model robustness and compute the RobustPass@k scores, which are then compared with the original pass@k scores. Following the previous study, we apply all perturbations spanning over four categories of format, syntax, and function to HumanEvalFix dataset, resulting in 1640, 2296 and 4264 samples in format, syntax, and function categories, respectively. Table 13 shows examples of these perturbations. Note that for each category, there are multiple perturbations that are applied by the library.

Table 12 An example of the HumanEvalFix benchmark for which the individual task-specific adapters are not able to solve the problem while the merged version of them using the weight-averaging method solves the problem successfully.

The input prompt consisting of the question, buggy code, test case, and the initial part of the answer that the model is expected to complete.	Question: Fix bugs in <code>largest_divisor</code> <pre>def largest_divisor(n: int) -> int: for i in reversed(range(n)): if n - i == 0: return i def check(largest_divisor): assert largest_divisor(3) == 1 assert largest_divisor(7) == 1 assert largest_divisor(10) == 5 assert largest_divisor(100) == 50 assert largest_divisor(49) == 7 check(largest_divisor)</pre> Answer: <pre>def largest_divisor(n: int) -> int:</pre>
The wrong code generated by each of the task-specific adapters of T1, T2, and T3. Note that all of these adapters have generated the same code for the given prompt.	<pre>def largest_divisor(n: int) -> int: for i in reversed(range(1, n + 1)): if n % i == 0: return i</pre>
The correct code generated by the merged version of the constituent adapters (T1, T2, and T3) derived by weight-averaging method.	<pre>def largest_divisor(n: int) -> int: for i in reversed(range(n)): if n % i == 0: return i return 1</pre>

Table 13 Different perturbation types on problem 15 of the HumanEvalFix dataset. Note that only the first part of the input prompt is included and the rest are eliminated as those are the same for all samples

non-perturbed code	Question: Fix bugs in largest_divisor. <pre>def largest_divisor(n: int) -> int: """For a given number n, find the largest number that divides n evenly, smaller than n >>> largest_divisor(15) 5 """ for i in reversed(range(n)): if n - i == 0: return i</pre>
perturbed by format (doc2comment)	Question: Fix bugs in largest_divisor. <pre>def largest_divisor(n: int) -> int: # For a given number n, find the largest number that divides n evenly, # smaller than n # >>> largest_divisor(15) # 5 for i in reversed(range(n)): if n - i == 0: return i</pre>
perturbed by syntax (DeadCodeInserter)	Question: Fix bugs in largest_divisor. <pre>def largest_divisor(n: int) -> int: """For a given number n, find the largest number that divides n evenly, smaller than n >>> largest_divisor(15) 5 """ for i in reversed(range(n)): while False: if n % i == 0: return i for i in reversed(range(n)): if n - i == 0: return i</pre>
perturbed by function (FuncRenameButterFinger)	Question: Fix bugs in larhest_divisor. <pre>def larhest_divisor(n: int) -> int: """ For a given number n, find the largest number that divides n evenly, smaller than n >>> larhest_divisor(15) 5 """ for i in reversed(range(n)): if n - i == 0: return i</pre>

For each of these categories, we apply the perturbations using 5 different seeds, following the methodology of the previous work Wang et al. (2023b). We then compute the RobustPass@ k scores with $k = 1$ for the individual task-specific adapters and two selected merged adapters. It is important to note that evaluating RobustPass@ k scores for all merged adapters is computationally intensive and not feasible given our available resources. Specifically, evaluating a single adapter across all four variants of the perturbed dataset takes at least 8–28x time, which is significantly more than the approximate 1 hour needed to compute the pass@ k score for a single model. As a result, we limit this experiment to a subset of adapters.

Table 14 reports the RobustPass@1 scores for individual task-specific adapters, along with results for merged adapter, T4-T2-T1, as this one showed the best performance when adapter order mattered (RQ3). When comparing the RobustPass@1 results of the merged adapter with those of the individual adapters, we do not observe significant differences. However, when comparing the RobustPass@1 scores with the original Pass@1 scores for

Table 14 RobustPass@1 scores of individual task-specific adapters vs. two merged adapters trained using StarCoder2 as the backbone model evaluated across four categories of perturbations.

	Tasks & Merging Methods		Format	Syntax	Function	Docstring
Individual Adapters	Development (T4)		13.01%	13.01%	03.66%	24.81%
	Improvement (T2)		13.11%	11.79%	04.47%	26.23%
	Program Repair (T1)		14.33%	13.11%	05.28%	29.01%
Merged Adapter (RQ3)	T4-T2-T1	weight-averaging	14.33%	12.80%	04.88%	29.01%
		ties	14.33%	12.50%	04.57%	29.01%
		dare-ties	14.13%	12.70%	04.88%	27.59%

the same adapters, we observe a general decline. These drops are significant across all categories, with the exception of the `function` category, which shows a substantial decrease.

The `function` category yields the lowest RobustPass@1 scores among all categories, highlighting the importance of having correct and meaningful function names in the input prompts to the model. It is important to note that in this category, function names appear in multiple parts of the input prompt, such as in the function definition and in the test cases. In our adapted code, we apply the same change all over and modify all occurrences of the function names. Among the remaining categories, `format` and `syntax` have similar scores, the latter being a bit lower. Among the different merging techniques, all three have similar scores, and there is no specific pattern showing either of them is more robust.

5.3 Continual merging vs. equal-weight merging

One of the main advantages of the continual merging approach is its memory efficiency compared to equal-weight merging. Specifically, the continual merging strategy requires storing only one merged adapter at a time along with the next adapter to be merged. Adapters are merged sequentially, eliminating the need to load and store all of them simultaneously. This property becomes especially important when dealing with a large number of task-specific adapters or when the adapter sizes are substantial. Notably, this memory efficiency is achieved without a significant drop in the performance of the merged adapters on the APR task. In fact, the continual merging approach can still yield high performance on the APR benchmark.

As mentioned earlier, when comparing the results of continual merging with equal-weight merging, we do not observe a significant difference in overall performance. Nevertheless, since the order in which task-specific adapters are merged affects the final performance, better results can be achieved by carefully selecting the merging sequence. Consider the Pass@1 scores of the merged adapters involving three task-specific adapters, as reported in Tables 4 and 8 for equal-weight and continual merging, respectively. Interestingly, for every adapter combination merged using the equal-weight approach, we can identify a variant of continual merging with the same base tasks that outperforms it. This trend also holds for merged adapters involving four task-specific adapters. These observations suggest that continual merging can be more effective, provided that an optimal order of merging is identified.

5.4 Fraction of sign difference in adapter weight parameters

To further analyze the results, we consider the internal adapter parameters of the models using the fraction of sign difference (Nguyen and Le 2024), i.e., the portion of parameters

having opposite signs, of the participating adapters in the merging process. Most existing merging methods take into account the sign of the parameters when merging different models or adapters. For instance, in the TIES method, only parameters with matching sign values are averaged in the final step. Therefore, analyzing the Fraction of Sign Difference (FSD) is crucial to understanding whether there is a relation between the performance of merged adapters and the similarity of their sign patterns (low or high FSD). A higher FSD value indicates that the adapters have more parameters with differing sign values, whereas a lower FSD suggests fewer such differences. Intuitively, adapters with a lower FSD can be considered more similar to each other, and thus their performance is expected to be more aligned compared to adapters with a higher FSD.

Figure 6 shows the FSD of individual task-specific adapters that were trained on top of StarCoder2 model compared to other adapters. By analyzing this figure with the performance of each adapter reported in Table 3, we observe that the tasks with a lower FSD, perform relatively closer to each other in terms of pass@k score. For example, T1 and T2 have the lowest FSD and their performance is also very close to each other being 28.66% and 27.90%, respectively. Similarly, T5 has the highest FSD with T1 and their performance difference is also larger than all the other tasks compared with T1. However, the same observation does not hold for T4 and T3 in comparison with T1. While T3 performs closer to T1, it has a higher FSD which shows that there is no general correlation between the FSD and the performance of the individual adapters.

On the other hand, having the FSD values between different task adapters, we cannot understand which adapters will perform better or worse compared to a specific adapter like T1. For instance, while T3 performs better than T1, it has a higher FSD compared to the FSD value of T1 and T2. In fact, the FSD can only be helpful to understand the relative similarity of task-adapters and therefore, we can conclude which adapters will lead to a smaller change after applying merging.

As it is observed from Fig. 6b, the same pattern exists for the Granite model. As mentioned before, the FSD values do not represent the performance drop or improvement of the adapters on T1. But, they represent the similarity of task-specific adapters, and therefore their performance scores' closeness to each other. As such, we observe that the FSD of T1 and T2 is still lower than the FSD of T1 and other adapters, although for this model, the performance of the T2 task is better than T1.

In summary, when having the FSD value of APR task relative to other task-specific adapters, we can decide which task-specific adapters will have a higher impact on the APR




Fig. 6 The Fraction of Sign Difference (FSD) between individual task-specific adapters trained on top of StarCoder2 and Granite models. The values show the percentage of parameters having different signs

task if they are merged together, in terms of having a closer performance to APR. In other words, when merging APR adapter with task-specific adapters having a high FSD value, the resulting adapter will perform significantly different (could be higher or lower) compared to the case where we merge APR with a task-specific adapter having a low FSD. This is supported by the differences in the performance that we observed in Table 4, specifically when on adapter is merged with APR (T1). There is however, no specific conclusion when more adapters are merged with APR, because it depends on other adapters that are added as third or fourth ones.

We also compute the FSD values between the merged adapters of each RQ and APR adapter. As the FSD differences of the adapters trained on StarCoder2 and its performance on all RQs are better than Granite, in the following, we only consider the StarCoder2 model. Figure 7 shows these values for RQ1 adapters and APR. The first observation is that the FSD value of adapters merged with ties method are all similar to the ones merged with weight-averaging method. This happens due to the fact that ties merging technique is similar to weight-averaging. In ties merging, after assigning the sign vectors, the final parameters are averaged, similar to weight-averaging.

Moreover, when comparing the values of merged adapters with different number of adapters, it can be observed that the adapters with lower constituent task-adapters have a lower FSD for ties and weight-averaging methods. As the number of adapters increases in these methods, the FSD values get higher with the last merged adapter having the highest FSD. Almost the opposite trend can be seen for the dare_ties method meaning that the merged adapters with more constituent adapters have relatively lower FSD values. Therefore, we cannot correlate the number of adapters in each merged adapter with their FSD values with respect to the APR adapter.

Additionally, the overall values of FSD for the adapters of RQ1 vs. APR adapter are lower than the FSD values of single task-specific adapters vs. APR adapter across all methods and merged adapters. This suggests that after merging APR with other similar adapters, the resulting adapters become more similar to APR in terms of their FSD.




Fig. 7 The Fraction of Sign Difference (FSD) between the merged adapters of RQ1 vs. APR adapter trained on StarCoder2 model

It is worth mentioning that such results do not correlate with the performance of the merged adapters. For example, a lower FSD value for the merged adapters does not necessarily show a high performance for that adapter. More specifically, this can be observed when considering the results of T1-T2 and T1-T3 tasks. Although T1-T2 has a lower FSD value, it performs worse than T1-T3 merged adapter. This happens across all merging methods.

Figure 8 reports the FSD values of the merged adapters in the second RQ again with respect to the APR adapter. As it is observed the overall values are higher than the values of the first RQ. This could be valid and expected as in this RQ, we eliminate the APR adapter from the merged adapter. Hence, the merged adapter of other task-specific adapters should have more parameters with different sign values compared with APR. This unveils the reason behind having a lower performance for the merged adapters of non-APR task-specific adapters compared to the first RQ where all the merged adapters include the APR adapter and their performance are higher in general.

Unlike the previous RQ, the FSD values of the merged adapters get lower as the number of the tasks involved in the merged adapter increases. This trend is consistent for all three merging methods.

Similar to previous observations, the low/high score of FSD is not related to the performance of the merged adapters. For example, the merged adapter T2-T3-T4-T5 has the lowest FSD in Fig. 8, but, also the a relatively lower performance compared to when two adapters are merged (See Table 6).

Figures 9 and 10 show the FSD results of merged adapters having 3 and 4 adapters, respectively. The first difference of the FSD results of RQ3 is that in continual merging approach, the FSD values of the ties and weight-averaging methods are no longer the same. This can be due to different weights considered in the continual merging, so the sign differences can vary.

Recall that in the third RQ, the position of the adapters in the merged adapters played a significant role. By having a closer look at these two figures, we notice that as we involve the APR adapter in the last step of the continual merging, the FSD values of the merged




Fig. 8 The Fraction of Sign Difference (FSD) between the merged adapters of RQ2 vs. APR adapter trained on StarCoder2 model




Fig. 9 The Fraction of Sign Difference (FSD) between the merged adapters of RQ3 having length of 3 vs. APR adapter trained on StarCoder2 model




Fig. 10 The Fraction of Sign Difference (FSD) between the merged adapters of RQ3 having length of 4 vs. APR adapter trained on StarCoder2 model

adapters vs. the APR adapter becomes lower. This is in agreement with the fact that the lastly added adapters in the merged adapters of RQ3 have the highest impact in the merged adapter. Similarly, the merged adapters that contain the APR adapter in their first steps of merging have a higher FSD value which further justifies the above mentioned fact.

The pass@k results of the third RQ showed that the performance of the merged adapters are highly dependent on the performance of the last adapters in the merged adapter. By analyzing these results and the FSD values for this RQ, the impact of the last adapters on

the merged adapters can be further observed. This observation is however different from our previous discussions in this section. For the other two RQs, and single task adapters, would could not observe a direct relation among the pass@k scores and the FSD values. But we can have a stronger conclusion for this relation in case of continual merging.

5.5 The state of the solved problems

To analyze the ratio of the problems that were solved by different merged adapters, we illustrate the number of unique problems that were solved by each adapter for T4-T2-T1 adapter as shown in Venn diagrams of Fig. 11. As it is observed, many problems are solved by any of the adapters, i.e., T4-T2-T1 from RQ3, T1-T2-T4 from RQ1 and APR adapter itself. However, there are some problems that can only be solved using merged adapters (either equal-weight averaging or continual merging); and the number of problems that are uniquely solved by merged adapters are more than the number of problems that are uniquely solved by




Fig. 11 Venn diagrams of the number of problems solved by two merged adapters obtained by equal-weight merging (RQ1) and continual merging (RQ3) approaches along with the APR adapter across three merging methods

solved by APR adapter. Different merging methods, i.e., dare-ties, ties, and weigh-averaging affect these numbers. For example, with dare-ties method, no problem is solved only with T1-T2-T4, but with the ties and weight-averaging methods, there are 3 and 2 problems that are solved only with these merged adapter (see plots related to the RQ1 of Fig. 11). Similarly, when considering T4-T2-T1 adapter, we can observe that 2 problems are solved using each of dare-ties and weight-averaging methods and no problem is with ties. But overall, there are 8, 8, and 11 problems that can only be solved with the merged methods using dare-ties, ties, and weight-averaging, respectively, compared to 2, 0, and 3 problems that can only be solved by APR adapter. This examination emphasizes our previous discussions about the benefits of using merged adapters for APR task.

5.6 Incorrect and correct problems when changing merging technique

We also calculated the number of problems that were correct using APR adapter and became incorrect using the merged adapter, and vice versa. These results are calculated both in RQ1 and RQ3, using the merged adapters T1-T2-T4 and T4-T2-T1, respectively. Note that the merged adapters are the same, and only their order matters in RQ3. Table 15 shows the results for pass@10 scores. The $A \rightarrow B$ in the table represents the number of problems that were correct using A , but were incorrect in B . For example, $APR \rightarrow \text{weight-averaging}$ value for equal-weight averaging (RQ1) being 5 means that there were five problems that were correctly solved using APR adapter, but are incorrect using weight-averaging . Two patterns can be seen from the values of this table. First, number of incorrect problems increases when we have a $\text{merged-model} \rightarrow APR$, compared to $APR \rightarrow \text{merged-model}$. This shows the benefits of merged models, even though being small. Second, when using continual merging on the weight-averaging merging technique, the number of incorrect problems decreases, compared to equal-weight averaging. This number from continual merging is on-par for other two merging techniques, ties and dare-ties compared to equal-weight averaging. It is worth mentioning that, the numbers are small, representing 2–5% of the 165 benchmark problems that are incorrectly solved when we change the models. Using pass@1 scores, we noticed a slightly different results, where the number of incorrect problems increases when we have $APR \rightarrow \text{merged-model}$, compared to $\text{merged-model} \rightarrow APR$. A similar trend is seen for continual merging vs. equal-weight averaging.

Table 15 Number of problems that are correctly solved by one model and become incorrect when solving with the other model, using pass@10 scores from StarCoder2 model for T1-T2-T4 and T4-T2-T1 merged adapters. The $A \rightarrow B$ in the table represents the number of problems that were correct using A , but were incorrect in B .

Models	# Incorrect problems equal-weight averaging (RQ1)	# Incorrect problems continual merging (RQ3)
$APR \rightarrow \text{weight-averaging}$	5	3
$\text{weight-averaging} \rightarrow APR$	6	8
$APR \rightarrow \text{ties}$	6	6
$\text{ties} \rightarrow APR$	8	9
$APR \rightarrow \text{dare-ties}$	4	4
$\text{dare-ties} \rightarrow APR$	9	5

Table 16 represents the number of correctly solved problems using APR and three merging methods, with equal-weight averaging compared to continual merging. The results are based on pass@10 scores for T1-T2-T4 and T4-T2-T1 merged adapters, RQ1 and RQ3, respectively. Using continual merging, the number of correctly solved problems increases slightly using weight-averaging and ties, and decreases for dare-ties. This result emphasize the benefit of continual merging, which also can depend on the merging approach.

5.7 Distribution of task-specific datasets

To better understand the underlying reasons for the performance of different task-adapters and their merged versions, we further analyze the distribution of the initial datasets used in this study. To this end, we use a small LLM, i.e., allMiniLM-L6-v2, to embed the five code datasets. This model produces fixed-size vectors of dimension 384 for each data point, representing its contextual meaning. After obtaining the embeddings, we reduce their dimensionality to 2D to visualize their distribution using t-SNE. We use APR as the base task for comparison and separately compare each other task-specific dataset with APR. Figure 12 shows these pair-wise data distribution comparisons for the five tasks used in this study.

As shown in the figure, the Improvement and Development tasks highly overlap with the APR task, while the Misc and Test & QA tasks show lower overlap with APR. This figure suggests that although the datasets form different clusters, they are relatively close in terms of overall semantic meaning. As a result, their individual performances are expected to be similar, which is consistent with our experimental results. Moreover, recall the performance of the individual task-adapters in Table 3. As reported, the Improvement task has the smallest absolute performance difference compared to APR across both models and evaluation metrics, with pass@1 differences of 0.76% and 0.79% for the StarCoder2 and Granite models, and pass@10 differences of 1.12% and 1.21% for the same models, respectively. This is expected, as the source datasets of these tasks have a high degree of overlap. Likewise, the Test & QA dataset shows the lowest similarity to the APR task, leading to the largest performance gap with APR task across both models. The Misc task, however, achieves the highest individual performance despite having lower similarity to the APR distribution, as shown in Fig. 12. This indicates that final model performance cannot be inferred solely from data distribution.

5.8 Comparison with other approaches

This study lies at the intersection of multiple domains and can be useful to researchers in these areas. First, we work with code datasets aimed at addressing software engineering problems. In particular, we focus on the APR task and four closely related tasks: Improvement, Development, Test & QA, and Misc. For these tasks, we select base code LLMs that

Table 16 Number of problems that are correctly solved by each model, using pass@10 scores of StarCoder2 model for T1-T2-T4 and T4-T2-T1 merged adapters.

Models	# Correct problems equal-weight averaging (RQ1)	# Correct problems continual merging (RQ3)
APR	67	67
weight-averaging	68	72
ties	69	70
dare-ties	72	68




Fig. 12 Pair-wise distribution of task-specific datasets used in this study, with the APR task as the base for comparison. As shown in the figures, the Improvement and Development tasks (top-left and top-right) highly overlap with the APR task, while the Misc and Test & QA tasks (bottom-left and bottom-right) show lower overlap with APR

cannot properly generate human-readable outputs and therefore require fine-tuning with instruction-following datasets, to align them with their target domain and enable human-readable responses (Muennighoff et al. 2023; Wang et al. 2023c). This alignment step is essential for creating usable specialized models from pre-trained LLMs (Wang et al. 2023c; Ouyang et al. 2022b). Our experiments use code LLMs specifically developed for coding tasks and software engineering problems (Nijkamp et al. 2023; Chen et al. 2021; Hu et al. 2020; Fu et al. 2022).

Model merging is a promoted technique that aims to combine two or more specialized models into a single multi-task model. This approach differs from studies that develop multi-task models by training a single model on multiple tasks simultaneously (Bai et al. 2009; Yuan et al. 2013; Luong et al. 2015). In model merging approach that we adopted, multiple specialized models are trained separately and then merged without further training. This idea is a relatively well-established area in the field of neural networks and deep learning (Izmailov et al. 2018; Singh and Jaggi 2020) and has recently shown promising results when applied to LLMs (Yadav et al. 2024; Yu et al. 2023; Wortsman et al. 2022; Ilharco et al. 2022). Since its introduction, this method has also been explored in other domains, such as computer vision (Chen et al. 2024; Stoica et al. 2024; Sung et al. 2022).

Most of the studies apply merging to the full parameter space of models. In contrast, our study focuses on merging *adapters* rather than full models. The main reason is that

PEFT methods provide a more efficient way to fine-tune LLMs in resource-constrained environments by updating only a small subset of model parameters during the fine-tuning phase (Poth et al. 2023; Ben Zaken et al. 2022; Li and Liang 2021; Liu et al. 2024; Shi et al. 2023). Since we aim to conduct a comprehensive set of experiments, we adopt a PEFT method, specifically LoRA (Hu et al. 2022), to fine-tune adapters and apply merging only to these LoRA adapters, without further training. This approach is time-efficient and remains relatively underexplored, with only a few studies focusing on merging adapters instead of full model parameters (Chronopoulou et al. 2023).

5.9 Hardware constraints

We outline our hardware constraints and provide additional details on the resources required to reproduce the experiments. The runtime of each fine-tuning experiment varies across tasks and models. Since we did not record exact runtimes during execution, the values reported here are estimates based on a small subset of experiments. Overall, fine-tuning a single task on each model takes approximately 12 hours. For inference, obtaining one pair of pass@1 and pass@10 scores for a single task on one model takes about 1 hour. In contrast, computing RobustPass@k scores requires substantially more time, ranging from approximately 8 to 28 hours depending on the perturbation category. The merging process is performed on CPU and takes only a few minutes per experiment. Except for the merging phase and other analytical components, all fine-tuning and inference experiments are conducted on a single NVIDIA Tesla V100 32GB GPU.

6 Implications and future research directions

6.1 Implications for researchers

Our study is the first to explore merging the learned parameters of models without additional training. We investigated the capabilities of merged models within the context of parameter-efficient fine-tuning approaches for automated program repair. Our results demonstrate that merging adapters from related software engineering tasks can benefit complex tasks like APR. Interestingly, the number of merged adapters or even the degree to which their parameter signs differ from those of the APR adapter, does not strongly correlate with performance. Instead, the effectiveness of the merged adapter depends more on the nature of the tasks being merged. Moreover, not all task-specific adapters are beneficial to a given target task. In our experiments, for instance, merging the T5 adapter negatively impacted the performance of the merged model. This highlights the importance of carefully selecting which task adapters to include in the merge. A promising direction for future work is to investigate which tasks contribute the most useful knowledge when merging adapters for a specific target task, and find the combination of best merging adapters.

In this study, we empirically explored the number and order of tasks for merging. However, exhaustively evaluating all possible task combinations and orders is nontrivial. A promising direction for future research is to develop automated techniques that can identify which tasks and in which order they should be merged to best benefit a given target task.

Additionally, our continual merging results suggest that the order in which adapters are merged affects the final performance, either positively or negatively. In this study, we explored different merging orders empirically. Based on the initial performance of individual adapters on the APR task, we were able to anticipate which tasks might be more beneficial and where they should be positioned in the merged model. Although we currently lack an automated method to identify the most beneficial tasks for a given target task, our findings indicate that placing the most effective adapter last in the continual merging sequence can improve performance. This is because the last adapter tends to have the greatest influence in the merged model, effectively assigning it a higher weight. For APR, this strategy can help ensure that the most relevant knowledge is prioritized, leading to better overall results.

In this study, we used an existing dataset comprising five tasks, including APR. This choice was particularly important, as we required instruction-tuning data compatible with the StarCoder2 and Granite models. After evaluating the results, we observed that the individual adapters performed similarly to APR on the APR benchmark (i.e., the HumanEval-Fix dataset). Upon manual inspection, we found that while the task-specific datasets were well-categorized for each task, the APR benchmark contained a more diverse set of samples, including instances from tasks beyond APR. This may explain why the APR-specific adapter did not outperform other task-specific adapters on the APR benchmark. Although this observation does not impact the main findings or conclusions of our study, we recommend that future research efforts focus on curating more clearly defined task-specific datasets and dedicated benchmark datasets to ensure accurate evaluation.

Moreover, based on the provided t-SNE plots, since the tasks we experimented with are closely related, their merging scores are also very similar. Researchers can investigate the impact of merging a broader variety of tasks for APR or other target tasks. For example, tasks such as comment generation or clone detection could be incorporated into the merged adapters to enhance performance on APR. More generally, researchers can curate new datasets and explore merging techniques across a wider range of software engineering tasks to better understand their effectiveness and transferability.

6.2 Implications for practitioners

Merging models without any additional training, as explored in this work, offers several practical benefits. Given the high cost of dataset curation and labeling, model merging enables practitioners to reuse existing knowledge from related tasks for new tasks. While we do not claim that merged models generalize across all software engineering tasks, our findings suggest that merged models exhibit generalization capabilities for automated program repair. Similar trends have been reported in the NLP domain, where merged models demonstrate promising performance across diverse tasks (Akiba et al. 2025).

This indicates that merging existing models may be a viable strategy worth exploring before investing in dataset creation for new SE tasks. The technique is particularly valuable when working with private or confidential codebases, as it allows knowledge transfer without requiring access to sensitive data for training, by leveraging the learned knowledge of other models or tasks. Beyond reusing pre-trained models, merging also contributes to energy savings, aligning with the goals of sustainable and environmentally responsible AI (Shi et al. 2024).

Another notable finding from our results is that while different merging methods yield varying scores for the merged adapters, the performance differences across these techniques were not substantial. This suggests that in similar settings, practitioners could opt for simpler and more computationally efficient methods, such as straightforward weight averaging, instead of more complex merging approaches.

In the continual merging setting, where the objective is to maintain a single merged adapter at a time, we observed that the most recently added adapters have a significant impact on the final performance. Therefore, to optimize the merged adapter's effectiveness, higher-performing adapters should be added later in the merging sequence. Our continual merging approach has two important implications for practitioners. First, it enables adapters to be merged sequentially as they become available over time, removing the need to train adapters for all tasks simultaneously. Second, it allows to weight the adapters differently, giving practitioners the ability to control their influence on the merged model based on the target task.

7 Threats to validity

Threats to **internal validity** refer to internal factors that might affect the reliability of the results. These factors in our study are related to model checkpoints and the dataset used. The model checkpoints that we selected are pre-trained without any further fine-tuning process such as instruction-tuning. The instruction-tuning process will align both models to generate more proper output code snippets. In all cases, we used the best hyper-parameters and configurations, as recommended by the model developers. As we used the open source code and APIs, we anticipate low threat related to this factor. The other potential threat concerns the dataset used in our experiments. Although the task type of the records in the original dataset has been obtained by prompting GPT-4 model, and might include incorrect labels, the same dataset is used for all experiments and models. Therefore, the results are affected in the same way by the used dataset. Please note that such noisy labels might exist in other datasets as well, and their effect needs to be evaluated in a separate study, which is out of the scope of our current study. Similarly, for testing the trained models, we used existing benchmark datasets to alleviate any threats related to the choice of datasets. Another threat can be related to the effect of model architecture and the model size. We did not conduct experiments on the effect of model architecture on the results, and we chose same size models, to reduce the related threats. It is worthy to note that our results are concerned with the merging capability of the trained adapters. To reduce threats related to this point, we chose two models in our study.

Threats to **external validity** relate to the generalizability of our findings. In this work, we conducted experiments only on automatic program repair, along with the other four tasks for Python language. Although the techniques used in the merged models/adapters can be adopted for other tasks and languages, the obtained results are limited to the used tasks and programming language and might not be applicable to other areas or programming languages.

Threats to **construct validity** refers to misalignment among the test and what needs to be measured. In our work, we intend to compare the results for APR task, which is evaluated using pass@k metric. We found this metric to be more reliable than others, as in the benchmark, the fixed codes generated by the models will be evaluated by test cases, rather than pure consideration of textual similarity. We also tested the differences of the obtained results using statistical tests and also reported effect size.

Lastly, we expect low threats to the **conclusion validity**. As the models are frozen and we kept all the models' configurations the same, the results are expected to be related to the merging techniques. We used different evaluation metrics, supported by statistical tests and additional analysis in the discussion section. It is worth noting the difference in the input format of the benchmark dataset when computing $\text{pass}@k$ and $\text{RobustPass}@k$ results. The bigcode-evaluation-harness library used for $\text{pass}@k$ eliminates the docstrings from the benchmark, while the $\text{RobustPass}@k$ library (Wang et al. 2023b) requires docstrings, as some of the perturbations are also applied on them. Therefore, the direct comparison among the scores might not be the best approach. That being said, the conclusion that perturbations affect the performance in a negative way are still valid.

8 Conclusion

In this study, we explored the effect of merging different task-specific adapters in Code LLMs on automatic program repair. The merged models and adapters have shown promising results in other fields including computer vision and natural language processing, and have improved the performance of the model for in-domain data. However, their effect on code-related tasks is not explored. Our study is the first to address this gap by conducting experiments on two merging paradigms, as well as generalizability of the merged adapters to the APR task. Our results suggest that merging task-specific adapters can improve the performance of the models on APR benchmark. Additionally, our proposed continual merging approach can be effective if the order of adapters that are merged together is determined efficiently. Such results and insights could shed light on the use of other code-related tasks to improve the performance of a new task. This work provides insights on how to learn from other tasks and develop adapters that are obtained from multiple adapters by different merging techniques. Our code and datasets are shared publicly. Future directions for this research could consider other tasks to combine for APR or study the effect of merging techniques for other SE tasks, as well as experimenting with other LLMs and industrial evaluations.

Appendix: RQ3 Results with T3 Included in the Merged Adapters

In this section, we add the results for RQ3 (continual merging) where T3 is included as one of the constituent tasks in the merged adapters. Please note that here we still consider merged adapters with up to four task-specific adapters, similar to RQ3. Since T5 has the lowest performance, we exclude it from the considered combinations. Therefore, we experiment with APR (T1), Improvement (T2), Misc (T3), and Development (T4) tasks. Tables 17 and 18 report the $\text{pass}@1$ scores of the merged adapters, while Tables 19 and 20 report the $\text{pass}@10$ scores.

Recall that T3 has the highest performance among all individual adapters. Considering merged adapters composed of three tasks in Table 8, for the $\text{pass}@1$ score, the best-performing merged adapter using the StarCoder2 model is built from T1, T2, and T3, with the order T2–T1–T3, achieving a +2.93% performance improvement (using the dare-ties method) over the APR adapter. For the $\text{pass}@1$ score using the Granite model, the best-performing merged adapter is built from T1, T3, and T4, with the order T1–T4–T3, leading to a +2.74% performance improvement with the dare-ties method. In both cases, T3, the

highest-performing adapter, is placed last in the merged adapter. Similarly, across other merged adapters, adding T3 in the final step results in relatively higher performance compared to other orderings.

For merged adapters composed of four tasks and the pass@1 score reported in Table 18, similar observations can be made. In particular, the best-performing adapters are T1–T4–T2–T3 and T4–T1–T2–T3, both achieving a +3.14% performance improvement over the APR adapter using the StarCoder2 model with the ties method. For the Granite model, the best-performing adapter is T1–T2–T4–T3, again with T3 placed in the final position and using the dare-ties method. These results are consistent with the findings of RQ3 reported earlier.

For the pass@10 scores reported in Tables 19 and 20, the best-performing merged adapter is T2–T1–T3 using the dare-ties method with the StarCoder2 model, leading to a +4.17% performance improvement over the APR adapter. For the Granite model, the highest performance improvement is achieved with the T4–T1–T3 merged adapter using the dare-ties method, resulting in a +2.28% pass@10 improvement. For merged adapters composed of four task-specific adapters, the best-performing ones are T3–T4–T2–T1 for the StarCoder2 model with the dare-ties method (yielding a +4.4% performance improvement), and T3–T1–T2–T4 as well as T1–T3–T2–T4 for the Granite model using the weight-averaging method, both achieving a +3.11% performance improvement. Notably, these latter cases do not follow the previously observed pattern in which placing the highest-performing adapter in the final step yields the best performance. However, we still observe that for StarCoder2 T3 is placed in last for the best performing model, and for Granite model, the adapters with T3 as last adapter, being T1–T2–T4–T3, T1–T4–T2–T3, T2–T1–T4–T3, T2–T4–T1–T3, T4–T1–T2–T3, and T4–T2–T1–T3, have positive improvements with effect size of medium to large. This shows that having the high performing adapters ordering in the end are still of high importance.

Table 17 pass@1 scores of the merged adapters using continual merging approach built from three task-specific adapters including T3. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-av eraging
T1,T2,T3	T1-T2-T3	30.58% (+1.92)	31.4% (+2.74)	31.31% (+2.65)	18.2% (+2.04)	17.59% (+1.43)	17.96% (+1.8)
		M***	L***	L***	M***	M***	S***
	T1-T3-T2	29.94% (+1.28)	30.15% (+1.49)	30.06% (+1.4)	17.44% (+1.28)	17.26% (+1.1)	17.74% (+1.58)
		S***	S***	M***	M***	L***	M***
	T2-T1-T3	31.59% (+2.93)	31.4% (+2.74)	31.31% (+2.65)	18.32% (+2.16)	17.59% (+1.43)	17.96% (+1.8)
		L***	M***	L***	L***	M***	M***
	T2-T3-T1	30.82% (+2.16)	30.91% (+2.25)	30.82% (+2.16)	16.55% (+0.39)	16.98% (+0.82)	17.04% (+0.88)
		L***	M***	M***	S**	S***	S**
	T3-T1-T2	29.91% (+1.25)	30.15% (+1.49)	30.06% (+1.4)	17.2% (+1.04)	17.26% (+1.1)	17.74% (+1.58)
		S***	M***	S**	M***	M***	M***
	T3-T2-T1	30.21% (+1.55)	30.91% (+2.25)	30.82% (+2.16)	16.8% (+0.64)	16.98% (+0.82)	17.04% (+0.88)
		M***	M***	M***	S*	S**	S**
T1,T3,T4	T1-T3-T4	28.29% (-0.37)	26.98% (-1.68)	27.84% (-0.82)	18.26% (+2.1)	<u>18.23% (+2.07)</u>	18.23% (+2.07)
		S*	S***	S*	L***	L***	M***
	T1-T4-T3	28.9% (+0.24)	28.78% (+0.12)	29.39% (+0.73)	18.9% (+2.74)	17.93% (+1.77)	18.32% (+2.16)
		S*	S*	S*	L***	L***	M***
	T3-T1-T4	27.26% (-1.4)	26.98% (-1.68)	27.84% (-0.82)	18.35% (+2.19)	<u>18.23% (+2.07)</u>	18.23% (+2.07)
		S***	M***	S*	L***	L***	M***
	T3-T4-T1	28.63% (-0.03)	29.42% (+0.76)	29.39% (+0.73)	17.53% (+1.37)	16.89% (+0.73)	17.23% (+1.07)
		S*	S*	S*	M***	M***	S**
	T4-T1-T3	29.12% (+0.46)	28.78% (+0.12)	29.39% (+0.73)	18.45% (+2.29)	17.93% (+1.77)	18.32% (+2.16)
		S*	S*	S*	L***	M***	M***
	T4-T3-T1	30.03% (+1.37)	29.42% (+0.76)	29.39% (+0.73)	17.38% (+1.22)	16.89% (+0.73)	17.23% (+1.07)
		S**	S**	S**	M***	S**	S*

Table 18 pass@1 scores of the merged adapters using continual merging approach built from four task-specific adapters including T3. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-averaging
T1,T2,T4	T1-T2-T4	27.32% (-1.34) S*	27.35% (-1.31) S**	27.5% (-1.16) S***	17.84% (+1.68) M***	17.47% (+1.31) M***	17.99% (+1.83) S***
	T1-T4-T2	28.29% (-0.37) S*	28.81% (+0.15) S*	28.78% (+0.12) S*	17.47% (+1.31) M***	17.32% (+1.16) S***	17.23% (+1.07) S***
	T2-T1-T4	28.17% (-0.49) S*	27.35% (-1.31) S*	27.5% (-1.16) S***	17.71% (+1.55) M***	17.47% (+1.31) S**	17.99% (+1.83) M***
	T2-T4-T1	29.21% (+0.55) S*	28.96% (+0.3) S*	28.78% (+0.12) S*	17.07% (+0.91) S*	16.55% (+0.39) S*	16.86% (+0.7) S*
	T4-T1-T2	28.11% (-0.55) S**	28.81% (+0.15) S*	28.78% (+0.12) S*	17.47% (+1.31) M***	17.32% (+1.16) S***	17.23% (+1.07) S***
	T4-T2-T1	29.3% (+0.64) S**	28.96% (+0.3) S**	28.78% (+0.12) S*	17.01% (+0.85) S***	16.55% (+0.39) S*	16.86% (+0.7) S**
	T1-T2-T5	25.73% (-2.93) L***	25.88% (-2.78) M***	25.37% (-3.29) L***	15.21% (-0.95) S***	14.24% (-1.92) M***	15.24% (-0.92) S***
	T1-T5-T2	27.9% (-0.76) S*	28.11% (-0.55) S*	27.59% (-1.07) S***	16.28% (+0.12) S*	15.61% (-0.55) S**	15.91% (-0.25) S*
	T2-T1-T5	25.52% (-3.14) L***	25.88% (-2.78) M***	25.37% (-3.29) L***	14.36% (-1.8) S***	14.24% (-1.92) M***	15.24% (-0.92) S***
	T2-T5-T1	27.68% (-0.98) S***	28.2% (-0.46) S*	27.41% (-1.25) M***	15.7% (-0.46) S*	15.58% (-0.58) S*	15.58% (-0.58) S*
	T5-T1-T2	27.84% (-0.82) S*	28.11% (-0.55) S*	27.59% (-1.07) S***	15.82% (-0.34) S*	15.61% (-0.55) S*	15.91% (-0.25) S*
	T5-T2-T1	28.11% (-0.55) S*	28.2% (-0.46) S*	27.41% (-1.25) S***	15.73% (-0.43) S*	15.58% (-0.58) S**	15.58% (-0.58) S*
T1,T4,T5	T1-T4-T5	25.43% (-3.23) L***	24.63% (-4.03) L***	25.43% (-3.23) L***	15.49% (-0.67) S*	15.34% (-0.82) S*	15.55% (-0.61) S**
	T1-T5-T4	26.07% (-2.59) L***	25.09% (-3.57) L***	25.95% (-2.71) L***	16.92% (+0.76) S***	16.55% (+0.39) S*	17.26% (+1.1) S***
	T4-T1-T5	24.73% (-3.93) L***	24.63% (-4.03) L***	25.43% (-3.23) L***	15.7% (-0.46) S*	15.34% (-0.82) S**	15.55% (-0.61) S*
	T4-T5-T1	27.07% (-1.59) S***	25.73% (-2.93) M***	26.46% (-2.2) M***	16.49% (+0.33) S**	16.07% (-0.09) S*	16.1% (-0.06) S*
	T5-T1-T4	26.28% (-2.38) L***	25.09% (-3.57) L***	25.95% (-2.71) L***	16.55% (+0.39) S***	16.55% (+0.39) S*	17.26% (+1.1) S*
	T5-T4-T1	26.8% (-1.86) S***	25.73% (-2.93) L***	26.46% (-2.2) M***	16.37% (+0.21) S*	16.07% (-0.09) S*	16.1% (-0.06) S*

Table 19 pass@10 scores of the merged adapters using continual merging approach built from three task-specific adapters including T3. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2 dare-ties	ties	weight-averaging	Granite dare-ties	ties	weight-averaging
T1,T2,T4,T5	T1-T2-T4-T5	25.03% (-3.63)	24.33% (-4.33)	25.73% (-2.93)	15.85% (-0.31)	15.4% (-0.76)	16.28% (+0.12)
		L***	L***	L***	S*	S*	S*
		25.21% (-3.45)	25.55% (-3.11)	25.67% (-2.99)	17.04% (+0.88)	16.43% (+0.27)	17.1% (+0.94)
T1,T2,T4,T5	T1-T2-T5-T4	L***	M***	L***	S***	S*	S*
		25.43% (-3.23)	25.49% (-3.17)	25.73% (-2.93)	14.66% (-1.5)	14.82% (-1.34)	15.82% (-0.34)
		L***	L***	L***	S**	M***	S*
T1,T2,T4,T5	T1-T4-T5-T2	26.31% (-2.35)	27.29% (-1.37)	27.16% (-1.5)	17.13% (+0.97)	16.13% (-0.03)	16.37% (+0.21)
		M***	S*	M***	S**	S*	S*
		25.85% (-2.81)	26.46% (-2.2)	26.43% (-2.23)	17.65% (+1.49)	16.95% (+0.79)	17.87% (+1.71)
T1,T2,T4,T5	T1-T5-T2-T4	L***	M***	L***	L***	S*	M***
		27.47% (-1.19)	27.96% (-0.7)	27.38% (-1.28)	17.56% (+1.4)	16.77% (+0.61)	17.01% (+0.85)
		S*	S*	M***	M***	S***	S***
T1,T2,T4,T5	T2-T1-T4-T5	25.46% (-3.2)	24.33% (-4.33)	25.73% (-2.93)	15.85% (-0.31)	15.4% (-0.76)	16.28% (+0.12)
		L***	L***	L***	S*	S***	S*
		25.76% (-2.9)	25.55% (-3.11)	25.67% (-2.99)	17.23% (+1.07)	16.43% (+0.27)	17.1% (+0.94)
T1,T2,T4,T5	T2-T1-T5-T4	L***	M***	L***	S***	S*	M***
		25.4% (-3.26)	25.3% (-3.36)	25.55% (-3.11)	14.7% (-1.46)	14.7% (-1.46)	15.43% (-0.73)
		L***	L***	L***	S***	M***	S*
T1,T2,T4,T5	T2-T4-T5-T1	27.04% (-1.62)	27.35% (-1.31)	27.13% (-1.53)	15.79% (-0.37)	15.46% (-0.7)	16.13% (-0.03)
		M***	S**	M***	S*	S*	S*
		26.77% (-1.89)	27.16% (-1.5)	26.52% (-2.14)	17.44% (+1.28)	17.26% (+1.1)	17.16% (+1.0)
T1,T2,T4,T5	T2-T5-T1-T4	M***	S*	M***	M***	S***	S***
		27.59% (-1.07)	27.2% (-1.46)	27.32% (-1.34)	16.62% (+0.46)	16.43% (+0.27)	16.8% (+0.64)
		S*	S*	S***	S**	S*	S*
T1,T2,T4,T5	T4-T1-T2-T5	25.73% (-2.93)	25.49% (-3.17)	25.73% (-2.93)	14.94% (-1.22)	14.82% (-1.34)	15.82% (-0.34)
		L***	M***	L***	S**	M***	S*
		27.01% (-1.65)	27.29% (-1.37)	27.16% (-1.5)	16.46% (+0.3)	16.13% (-0.03)	16.37% (+0.21)
T1,T2,T4,T5	T4-T1-T5-T2	M***	S***	S***	S**	S*	S*
		25.4% (-3.26)	25.3% (-3.36)	25.55% (-3.11)	15.37% (-0.79)	14.7% (-1.46)	15.43% (-0.73)
		L***	L***	L***	S*	M***	S*
T1,T2,T4,T5	T4-T2-T1-T5	28.02% (-0.64)	27.35% (-1.31)	27.13% (-1.53)	15.98% (-0.18)	15.46% (-0.7)	16.13% (-0.03)
		S*	S*	M***	S*	S**	S*
		28.69% (+0.03)	28.38% (-0.28)	28.26% (-0.4)	16.89% (+0.73)	16.34% (+0.18)	16.86% (+0.7)
T1,T2,T4,T5	T4-T5-T1-T2	S*	S*	S*	S***	S*	S**
		28.57% (-0.09)	28.69% (+0.03)	29.02% (+0.36)	15.95% (-0.21)	16.13% (-0.03)	16.62% (+0.46)
		S*	S*	S*	S*	S*	S*
T1,T2,T4,T5	T5-T1-T2-T4	27.26% (-1.4)	26.46% (-2.2)	26.43% (-2.23)	17.87% (+1.71)	16.95% (+0.79)	17.87% (+1.71)
		S***	M***	L***	L***	S*	M***
		28.08% (-0.58)	27.96% (-0.7)	27.38% (-1.28)	17.5% (+1.34)	16.77% (+0.61)	17.01% (+0.85)
T1,T2,T4,T5	T5-T1-T4-T2	S*	S*	S*	M***	S**	S*
		27.41% (-1.25)	27.16% (-1.5)	26.52% (-2.14)	17.68% (+1.52)	17.26% (+1.1)	17.16% (+1.0)
		S***	S**	L***	M***	S***	S***
T1,T2,T4,T5	T5-T2-T1-T4	27.62% (-1.04)	27.2% (-1.46)	27.32% (-1.34)	16.52% (+0.36)	16.43% (+0.27)	16.8% (+0.64)
		S**	S**	S**	S*	S*	S*
		28.32% (-0.34)	28.38% (-0.28)	28.26% (-0.4)	16.74% (+0.58)	16.34% (+0.18)	16.86% (+0.7)
T1,T2,T4,T5	T5-T4-T1-T2	S*	S*	S*	S***	S*	S**
		28.41% (-0.25)	28.69% (+0.03)	29.02% (+0.36)	16.01% (-0.15)	16.13% (-0.03)	16.62% (+0.46)
		S*	S**	S*	S*	S*	S**

Table 20 pass@10 scores of the merged adapters using continual merging approach built from four task-specific adapters including T3. The results are reported across three merging methods for StarCoder2 and Granite models. Note that the order of the tasks being merged is important in this RQ as it affects the weight of each adapter used in the merged adapters. The rightmost adapters have higher weights. The tasks are listed as T1 (APR), T2 (Improvement), T3 (Misc), T4 (Development), and T5 (Test & QA)

Tasks	Order	StarCoder2			Granite		
		dare-ties	ties	weight-averaging	dare-ties	ties	weight-averaging
T1-T2,T3,T4	T1-T2-T3-T4	39.08% (-0.03)	39.72% (+0.61)	38.09% (-1.02)	21.66% (+1.82)	21.01% (+1.17)	22.6% (+2.76)
		S*	S*	S*	S***	S***	L***
		40.92% (+1.81)	40.92% (+1.81)	41.87% (+2.76)	22.68% (+2.84)	20.71% (+0.87)	21.62% (+1.78)
T1-T2-T3,T4	T1-T2-T3-T4	S**	S***	L***	L***	S**	M***
		38.18% (-0.93)	38.16% (-0.95)	39.49% (+0.38)	22.18% (+2.34)	21.47% (+1.63)	22.95% (+3.11)
		S**	S*	S*	M***	M***	L***
T1-T3-T4,T2	T1-T3-T4-T2	40.88% (+1.77)	40.51% (+1.4)	39.35% (+0.24)	22.28% (+2.44)	21.87% (+2.03)	22.17% (+2.33)
		S**	M***	S*	L***	M***	M***
		42.09% (+2.98)	41.73% (+2.62)	40.57% (+1.46)	21.94% (+2.1)	20.71% (+0.87)	21.82% (+1.98)
T1-T4-T2,T3	T1-T4-T2-T3	M***	M***	S**	M***	S*	L***
		39.62% (+0.51)	40.04% (+0.93)	40.04% (+0.93)	21.72% (+1.88)	21.0% (+1.16)	21.82% (+1.98)
		S*	S*	S**	M***	S***	M***
T2-T1-T3,T4	T2-T1-T3-T4	37.23% (-1.88)	39.72% (+0.61)	38.09% (-1.02)	22.12% (+2.28)	21.01% (+1.17)	22.6% (+2.76)
		M***	S*	S**	M***	M***	L***
		41.39% (+2.28)	40.92% (+1.81)	41.87% (+2.76)	22.02% (+2.18)	20.71% (+0.87)	21.62% (+1.78)
T2-T1-T4,T3	T2-T1-T4-T3	S**	S***	M***	M***	S***	M***
		38.29% (-0.82)	39.42% (+0.31)	37.22% (-1.89)	22.37% (+2.53)	21.77% (+1.93)	21.28% (+1.44)
		S*	S*	S***	M***	M***	S***
T2-T3-T1,T4	T2-T3-T1-T4	40.37% (+1.26)	41.22% (+2.11)	41.05% (+1.94)	21.75% (+1.91)	20.52% (+0.68)	21.95% (+2.11)
		S**	M***	M***	M***	S***	M***
		38.84% (-0.27)	42.36% (+3.25)	41.46% (+2.35)	22.12% (+2.28)	20.56% (+0.72)	21.35% (+1.51)
T2-T4-T1,T3	T2-T4-T1-T3	S*	L***	M***	M***	S*	S***
		42.68% (+3.57)	42.15% (+3.04)	41.97% (+2.86)	21.19% (+1.35)	21.31% (+1.47)	21.14% (+1.3)
		L***	L***	L***	S***	M***	S***
T3-T1-T2,T4	T3-T1-T2-T4	37.99% (-1.12)	38.16% (-0.95)	39.49% (+0.38)	22.57% (+2.73)	21.47% (+1.63)	22.95% (+3.11)
		S**	S*	S*	M***	M***	L***
		39.65% (+0.54)	40.51% (+1.4)	39.35% (+0.24)	21.64% (+1.8)	21.87% (+2.03)	22.17% (+2.33)
T3-T1-T4,T2	T3-T1-T4-T2	S*	S*	S*	M***	L***	L***
		37.59% (-1.52)	39.42% (+0.31)	37.22% (-1.89)	22.49% (+2.65)	21.77% (+1.93)	21.28% (+1.44)
		M***	S*	S***	M***	M***	M***
T3-T2-T1,T4	T3-T2-T1-T4	38.52% (-0.59)	41.22% (+2.11)	41.05% (+1.94)	21.45% (+1.61)	20.52% (+0.68)	21.95% (+2.11)
		S***	M***	M***	M***	S**	M***
		40.85% (+1.74)	40.0% (+0.89)	40.93% (+1.82)	22.64% (+2.8)	20.37% (+0.53)	21.75% (+1.91)
T3-T4-T1,T2	T3-T4-T1-T2	S***	S**	M***	L***	S*	M***
		43.51% (+4.4)	41.37% (+2.26)	41.88% (+2.77)	21.58% (+1.74)	20.25% (+0.41)	21.23% (+1.39)
		L***	M***	L***	M***	S*	M***
T4-T1-T2,T3	T4-T1-T2-T3	41.69% (+2.58)	41.73% (+2.62)	40.57% (+1.46)	21.97% (+2.13)	20.71% (+0.87)	21.82% (+1.98)
		M***	M***	S**	L***	M***	M***
		40.31% (+1.2)	40.04% (+0.93)	40.04% (+0.93)	22.5% (+2.66)	21.0% (+1.16)	21.82% (+1.98)
T4-T1-T3,T2	T4-T1-T3-T2	S*	S*	S*	L***	S**	L***
		41.67% (+2.56)	42.36% (+3.25)	41.46% (+2.35)	22.06% (+2.22)	20.56% (+0.72)	21.35% (+1.51)
		M***	L***	M***	M***	S**	M***
T4-T2-T1,T3	T4-T2-T1-T3	40.76% (+1.65)	42.15% (+3.04)	41.97% (+2.86)	20.84% (+1.0)	21.31% (+1.47)	21.14% (+1.3)
		S**	L***	M***	S**	M***	S***
		41.0% (+1.89)	40.0% (+0.89)	40.93% (+1.82)	21.41% (+1.57)	20.37% (+0.53)	21.75% (+1.91)
T4-T3-T1,T2	T4-T3-T1-T2	M***	S**	M***	M***	S*	M***
		41.42% (+2.31)	41.37% (+2.26)	41.88% (+2.77)	20.86% (+1.02)	20.25% (+0.41)	21.23% (+1.39)
		M***	M***	M***	S***	S*	S***

Author Contributions Meghdad Dehghan contributed to the methodology, conducted the experiments, performed the analysis, and wrote the original draft. Jie JW Wu, Fatemeh H. Fard, and Ali Ouni contributed to the methodology and supervised the project. All authors contributed to reviewing and editing the manuscript.

Funding This research is supported by a grant from the Natural Sciences and Engineering Research Council of Canada RGPIN-2019-05175, as well as support from computational resources and services provided by Advanced Research Computing at the University of British Columbia.

Data Availability We include all scripts used to obtain the results in our [GitHub repository](#). Additionally, we have uploaded the fine-tuned LoRA adapters to Hugging Face, and they are available for immediate download and use [here](#).

Declarations

Ethical Approval Not applicable.

Informed Consent Not applicable.

Conflicts of Interest The authors declare that they have no conflict of interest.

Clinical Trial Number Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ainsworth S, Hayase J, Srinivasa S (2022) Git re-basin: merging models modulo permutation symmetries. The Eleventh international conference on learning representations
- Akiba T, Shing M, Tang Y, Sun Q, Ha D (2025) Evolutionary optimization of model merging recipes. *Nature Machine Intelligence* 1–10
- Bai J, Zhou K, Xue G, Zha H, Sun G, Tseng B, Zheng Z, Chang Y (2009) Multi-task learning for learning to rank in web search. In: *Proceedings of the 18th ACM conference on Information and knowledge management*, pp 1549–1552
- Bavarian M, Jun H, Tezak N, Schulman J, McLeavey C, Tworek J, Chen M (2022) Efficient training of language models to fill in the middle. *arXiv preprint [arXiv:2207.14255](https://arxiv.org/abs/2207.14255)*
- Beeching E, Fourrier C, Habib N, Han S, Lambert N, Rajani N, Sanseviero O, Tunstall L, Wolf T (2023) Open LLM leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard
- Ben Zaken E, Goldberg Y, Ravfogel S (2022) BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In: Muresan S, Nakov P, Villavicencio A (eds) *Proceedings of the 60th annual meeting of the association for computational linguistics (Volume 2: Short Papers)*, Association for Computational Linguistics, Dublin, Ireland, pp 1–9. <https://doi.org/10.18653/v1/2022.acl-short.1>
- Chen Z, Kommrusch S, Tufano M, Pouchet LN, Poshyvanyk D, Monperrus M (2019) Sequencer: sequence-to-sequence learning for end-to-end program repair. *IEEE Trans Softw Eng* 47(9):1943–1959
- Chen M, Tworek J, Jun H, Yuan Q, Pinto HPdO, Kaplan J, Edwards H, Burda Y, Joseph N, Brockman G, et al. (2021) Evaluating large language models trained on code. *arXiv preprint [arXiv:2107.03374](https://arxiv.org/abs/2107.03374)*
- Chen H, Tao R, Zhang H, Wang Y, Li X, Ye W, Wang J, Hu G, Savvides M (2024) Conv-adapter: exploring parameter efficient transfer learning for convnets. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 1551–1561
- Chronopoulou A, Peters M, Fraser A, Dodge J (2023) AdapterSoup: weight averaging to improve generalization of pretrained language models. In: Vlachos A, Augenstein I (eds) *Findings of the association for computational linguistics: EACL 2023*, Association for Computational Linguistics, Dubrovnik, Croatia, pp 2054–2066 <https://doi.org/10.18653/v1/2023.findings-eacl.153>
- Cliff N (1993) Dominance statistics: ordinal analyses to answer ordinal questions. *Psychol Bull* 114(3):494–509
- Ding N, Qin Y, Yang G, Wei F, Yang Z, Su Y, Hu S, Chen Y, Chan CM, Chen W et al (2023) Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nat Mach Intell* 5(3):220–235
- Li R et al (2023) Starcoder: may the source be with you! *Transactions on Machine Learning Research*. <https://openreview.net/forum?id=KoFOg4lhaE>, reproducibility Certification
- Fu M, Tantithamthavorn C, Le T, Nguyen V, Phung D (2022) Vulrepair: a t5-based automated software vulnerability repair. In: *Proceedings of the 30th ACM joint European software engineering conference and symposium on the foundations of software engineering*, pp 935–947

- Goddard C, Siriwardhana S, Ehghaghi M, Meyers L, Karpukhin V, Benedict B, McQuade M, Solawetz J (2024) Arcee's mergekit: a toolkit for merging large language models. arXiv preprint [arXiv:2403.13257](https://arxiv.org/abs/2403.13257)
- Goues L, Dewey-Vogt C, Forrest M, Weimer S, W (2012) A systematic study of automated program repair: Fixing 55 out of 105 bugs for \$8 each. In: 34th International conference on software engineering (ICSE). IEEE, pp 3–13
- Haque MZ, Afrin S, Mastropaolo A (2025) A systematic literature review of parameter-efficient fine-tuning for large code models. [arXiv:2504.21569](https://arxiv.org/abs/2504.21569) arXiv preprint
- Houlsby N, Giurgiu A, Jastrzebski S, Morrone B, De Laroussilhe Q, Gesmundo A, Attariyan M, Gelly S (2019a) Parameter-efficient transfer learning for nlp. In: International conference on machine learning, PMLR, pp 2790–2799
- Houlsby N, Giurgiu A, Jastrzebski S, Morrone B, De Laroussilhe Q, Gesmundo A, Attariyan M, Gelly S (2019b) Parameter-efficient transfer learning for nlp. In: International conference on machine learning, PMLR, pp 2790–2799
- Howard J, Ruder S (2018) Universal language model fine-tuning for text classification. In: Gurevych I, Miyao Y (eds) Proceedings of the 56th annual meeting of the association for computational linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Melbourne, Australia, pp 328–339. <https://doi.org/10.18653/v1/P18-1031>
- Hu X, Li G, Xia X, Lo D, Jin Z (2020) Deep code comment generation with hybrid lexical and syntactical information. *Empir Softw Eng* 25:2179–2217
- Hu EJ, Shen Y, Wallis P, Allen-Zhu Z, Li Y, Wang S, Wang L, Chen W (2022) LoRA: low-rank adaptation of large language models. In: International conference on learning representations. <https://openreview.net/forum?id=nZeVKeeFYf9>
- Huang C, Liu Q, Lin BY, Pang T, Du C, Lin M (2023) LoraHub: efficient cross-task generalization via dynamic lora composition. arXiv preprint [arXiv:2307.13269](https://arxiv.org/abs/2307.13269)
- Huang C, Ye P, Chen T, He T, Yue X, Ouyang W (2024) Emr-merging: tuning-free high-performance model merging. *Adv Neural Inf Process Syst* 37:122741–122769
- Ilharco G, Ribeiro MT, Wortsman M, Gururangan S, Schmidt L, Hajishirzi H, Farhadi A (2022) Editing models with task arithmetic. [arXiv:2212.04089](https://arxiv.org/abs/2212.04089) arXiv preprint
- Izmailov P, Podoprikin D, Garipov T, Vetrov D, Wilson AG (2018) Averaging weights leads to wider optima and better generalization. [arXiv:1803.05407](https://arxiv.org/abs/1803.05407) arXiv preprint
- Jobstmann B, Griesmayer A, Bloem R (2005) Program repair as a game. In: Computer aided verification: 17th international conference, CAV 2005, Edinburgh, Scotland, UK, July 6–10, 2005. Proceedings 17, Springer, pp 226–238
- Kim TK (2015) T test as a parametric statistic. *Korean J Anesthesiol* 68(6):540–546
- Koyuncu A, Liu K, Bissyandé TF, Kim D, Klein J, Monperrus M, Le Traon Y (2020) Fixminer: mining relevant fix patterns for automated program repair. *Empir Softw Eng* 25:1980–2024
- Li XL, Liang P (2021) Prefix-tuning: optimizing continuous prompts for generation. In: Zong C, Xia F, Li W, Navigli R (eds) Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (Volume 1: Long Papers), Association for Computational Linguistics, Online, pp 4582–4597. <https://doi.org/10.18653/v1/2021.acl-long.353>
- Li Y, Mehrish A, Bhardwaj R, Majumder N, Cheng B, Zhao S, Zadeh A, Mihalcea R, Poria S (2023) Evaluating parameter-efficient transfer learning approaches on sure benchmark for speech understanding. In: ICASSP 2023–2023 IEEE international conference on acoustics. Speech and Signal Processing (ICASSP), IEEE, pp 1–5
- Liu X, Gao J, He X, Deng L, Duh K, Wang Yy (2015) Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In: Mihalcea R, Chai J, Sarkar A (eds) Proceedings of the 2015 conference of the North American chapter of the association for computational linguistics: human language technologies, Association for Computational Linguistics, Denver, Colorado, pp 912–921. <https://doi.org/10.3115/v1/N15-1092>
- Liu S, Gao C, Chen S, Nie LY, Liu Y (2020) Atom: commit message generation based on abstract syntax tree and hybrid ranking. *IEEE Trans Softw Eng* 48(5):1800–1817
- Liu K, Li L, Koyuncu A, Kim D, Liu Z, Klein J, Bissyandé TF (2021) A critical review on the evaluation of automated program repair systems. *J Syst Softw* 171:110817
- Liu H, Tam D, Muqeeth M, Mohta J, Huang T, Bansal M, Raffel CA (2022) Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Adv Neural Inf Process Syst* 35:1950–1965
- Liu S, Keung J, Yang Z, Liu F, Zhou Q, Liao Y (2024) Delving into parameter-efficient fine-tuning in code change learning: an empirical study. arXiv preprint [arXiv:2402.06247](https://arxiv.org/abs/2402.06247)
- Lozhkov A, Li R, Allal LB, Cassano F, Lamy-Poirier J, Tazi N, Tang A, Pykhtar D, Liu J, Wei Y, et al. (2024) Starcoder 2 and the stack v2: the next generation. arXiv preprint [arXiv:2402.19173](https://arxiv.org/abs/2402.19173)

- Luong MT, Le QV, Sutskever I, Vinyals O, Kaiser L (2015) Multi-task sequence to sequence learning. [arXiv:1511.06114](https://arxiv.org/abs/1511.06114) arXiv preprint
- Lutellier T, Pham HV, Pang L, Li Y, Wei M, Tan L (2020) Coconut: combining context-aware neural translation models using ensemble for program repair. In: Proceedings of the 29th ACM SIGSOFT international symposium on software testing and analysis. pp 101–114
- Männistö J, Attieh J, Tiedemann J (2025) A comparative study of PEFT methods for python code generation. In: Johansson R, Stymne S (eds) Proceedings of the joint 25th nordic conference on computational linguistics and 11th baltic conference on human language technologies (NoDaLiDa/Baltic-HLT 2025), University of Tartu Library, Tallinn, Estonia, pp 390–396. <https://aclanthology.org/2025.nodalida-1.42/>
- Matena MS, Raffel CA (2022) Merging models with fisher-weighted averaging. *Adv Neural Inf Process Syst* 35:17703–17716
- Mesbah A, Rice A, Johnston E, Glorioso N, Aftandilian E (2019) Deepdelta: learning to repair compilation errors. In: Proceedings of the 2019 27th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering, pp 925–936
- Mishra M, Stallone M, Zhang G, Shen Y, Prasad A, Soria AM, Merler M, Selvam P, Surendran S, Singh S (2024) Granite code models: a family of open foundation models for code intelligence. [arXiv:2405.04324](https://arxiv.org/abs/2405.04324) arXiv preprint
- Motwani M, Brun Y (2020) Automatically repairing programs using both tests and bug reports. [arXiv:2011.08340](https://arxiv.org/abs/2011.08340) arXiv preprint
- Muennighoff N, Liu Q, Zebaze AR, Zheng Q, Hui B, Zhuo TY, Singh S, Tang X, Werra V, Longpre L, S (2023) Octopack: instruction tuning code large language models. In: The twelfth international conference on learning representations
- Nguyen T, Le T (2024) Generalizability of mixture of domain-specific adapters from the lens of signed weight directions and its application to effective model pruning. [arXiv:2402.10639](https://arxiv.org/abs/2402.10639) arXiv preprint
- Nguyen H, Qi D, Roychoudhury A, Chandra S (2013) Semfix: program repair via semantic analysis. In: 35th International conference on software engineering (ICSE). IEEE, pp 772–781
- Nijkamp E, Pang B, Hayashi H, Tu L, Wang H, Zhou Y, Savarese S, Xiong C (2023) Codegen: an open large language model for code with multi-turn program synthesis. In: The eleventh international conference on learning representations. https://openreview.net/forum?id=iaYcJKpY2B_
- Otake S, Kawakami R, Inoue N (2023) Parameter efficient transfer learning for various speech processing tasks. In: ICASSP 2023–2023 IEEE international conference on acoustics, speech and signal processing (ICASSP), IEEE, pp 1–5
- Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, Zhang C, Agarwal S, Slama K, Ray A et al (2022a) Training language models to follow instructions with human feedback. *Adv Neural Inf Process Syst* 35:27730–27744
- Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, Zhang C, Agarwal S, Slama K, Ray A et al (2022b) Training language models to follow instructions with human feedback. *Adv Neural Inf Process Syst* 35:27730–27744
- Pandey SK, Mishra RB, Tripathi AK (2020) Bpdet: an effective software bug prediction model using deep representation and ensemble learning techniques. *Expert Syst Appl* 144:113085
- Papineni K, Roukos S, Ward T, Zhu WJ (2002) Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting of the association for computational linguistics, pp 311–318
- Pfeiffer J, Kamath A, Rücklé A, Cho K, Gurevych I (2021) AdapterFusion: non-destructive task composition for transfer learning. In: Merlo P, Tiedemann J, Tsarfaty R (eds) Proceedings of the 16th conference of the European chapter of the association for computational linguistics: main volume, Association for Computational Linguistics, Online, pp 487–503. <https://doi.org/10.18653/v1/2021.eacl-main.39>
- Poth C, Sterz H, Paul I, Purkayastha S, Engländer L, Imhof T, Vulić I, Ruder S, Gurevych I, Pfeiffer J (2023) Adapters: a unified library for parameter-efficient and modular transfer learning. [arXiv:2311.11077](https://arxiv.org/abs/2311.11077) arXiv preprint
- Prabhakar A, Li Y, Narasimhan K, Kakade S, Malach E, Jelassi S (2024) Lora soups: merging loras for practical skill composition tasks. [arXiv:2410.13025](https://arxiv.org/abs/2410.13025) arXiv preprint
- Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, Zhou Y, Li W, Liu PJ (2020) Exploring the limits of transfer learning with a unified text-to-text transformer. *J Mach Learn Res* 21(140):1–67
- Ren S, Guo D, Lu S, Zhou L, Liu S, Tang D, Sundaresan N, Zhou M, Blanco A, Ma S (2020) Codebleu: a method for automatic evaluation of code synthesis. [arXiv preprint arXiv:2009.10297](https://arxiv.org/abs/2009.10297)
- Shi E, Wang Y, Zhang H, Du L, Han S, Zhang D, Sun H (2023) Towards efficient fine-tuning of pre-trained code models: an experimental study and beyond. In: Proceedings of the 32nd ACM SIGSOFT international symposium on software testing and analysis, pp 39–51
- Shi J, Yang Z, Lo D (2024) Efficient and green large language models for software engineering: vision and the road ahead. *ACM Transactions on Software Engineering and Methodology*

- Silva A, Fang S, Monperrus M (2023) Repairllama: efficient representations and fine-tuned adapters for program repair. [arXiv:2312.15698](https://arxiv.org/abs/2312.15698) arXiv preprint
- Singh SP, Jaggi M (2020) Model fusion via optimal transport. *Adv Neural Inf Process Syst* 33:22045–22055
- Stoica G, Bolya D, Björner JB, Ramesh P, Hearn T, Hoffman J (2024) Zipit! merging models from different tasks without training. In: The twelfth international conference on learning representations. <https://openreview.net/forum?id=LEYUkvduhQ>
- Sung YL, Cho J, Bansal M (2022) V1-adapter: parameter-efficient transfer learning for vision-and-language tasks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 5227–5237
- Torrey L, Shavlik J (2010) Transfer learning. *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*. IGI global, pp 242–264
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. *Advances in Neural Information Processing Systems* 30
- Wang Y, Agarwal S, Mukherjee S, Liu X, Gao J, Awadallah AH, Gao J (2022) AdaMix: Mixture-of-adaptations for parameter-efficient model tuning. In: Goldberg Y, Kozareva Z, Zhang Y (eds) *Proceedings of the 2022 conference on empirical methods in natural language processing*, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, pp 5744–5760. <https://doi.org/10.18653/v1/2022.emnlp-main.388>
- Wang D, Chen B, Li S, Luo W, Peng S, Dong W, Liao X (2023a) One adapter for all programming languages? Adapter tuning for code search and summarization. In: *2023 IEEE/ACM 45th international conference on software engineering (ICSE)*, IEEE, pp 5–16
- Wang S, Li Z, Qian H, Yang C, Wang Z, Shang M, Kumar V, Tan S, Ray B, Bhatia P, Nallapati R, Ramanathan MK, Roth D, Xiang B (2023b) ReCode: robustness evaluation of code generation models. In: Rogers A, Boyd-Graber J, Okazaki N (eds) *Proceedings of the 61st annual meeting of the association for computational linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Toronto, Canada, pp 13818–13843. <https://doi.org/10.18653/v1/2023.acl-long.773>
- Wang Y, Kordi Y, Mishra S, Liu A, Smith NA, Khashabi D, Hajishirzi H (2023c) Self-instruct: Aligning language models with self-generated instructions. In: *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: long papers)*, pp 13484–13508
- Wei Y, Pei Y, Furia CA, Silva LS, Buchholz S, Meyer B, Zeller A (2010) Automated fixing of programs with contracts. In: *Proceedings of the 19th international symposium on Software testing and analysis*, pp 61–72
- Weimer W, Nguyen T, Goues L, Forrest C, S (2009) Automatically finding patches using genetic programming. In: *IEEE 31st international conference on software engineering*. IEEE, pp 364–374
- Weimer W, Forrest S, Le Goues C, Nguyen T (2010) Automatic program repair with evolutionary computation. *Commun ACM* 53(5):109–116
- Weyssow M, Zhou X, Kim K, Lo D, Sahraoui H (2023a) Exploring parameter-efficient fine-tuning techniques for code generation with large language models. *arXiv preprint arXiv:2308.10462*
- Weyssow M, Zhou X, Kim K, Lo D, Sahraoui H (2023b) 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*, pp 1470–1482
- Wong CP, Santiesteban P, Kästner C, Goues L, C (2021) Varfix: balancing edit expressiveness and search effectiveness in automated program repair. In: *Proceedings of the 29th ACM joint meeting on European software engineering conference and symposium on the foundations of software engineering*, pp 354–366
- Wortsman M, Ilharco G, Gadre SY, Roelofs R, Gontijo-Lopes R, Morcos AS, Namkoong H, Farhadi A, Carmon Y, Kornblith S (2022) Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: *International conference on machine learning*. PMLR, pp 23965–23998
- Wu X, Huang S, Wei F (2024) Mixture of lora experts. [arXiv:2404.13628](https://arxiv.org/abs/2404.13628) arXiv preprint
- Xia CS, Wei Y, Zhang L (2023) Automated program repair in the era of large pre-trained language models. In: *Proceedings of the 45th international conference on software engineering*, IEEE Press, ICSE '23, pp 1482–1494. <https://doi.org/10.1109/ICSE48619.2023.00129>
- Xin Y, Luo S, Liu X, Zhou H, Cheng X, Lee CE, Du J, Wang H, Chen M, Liu T et al (2024) V-petl bench: a unified visual parameter-efficient transfer learning benchmark. *Adv Neural Inf Process Syst* 37:80522–80535
- Xuan J, Martinez M, Demarco F, Clement M, Marcote SL, Durieux T, Le Berre D, Monperrus M (2016) Nopol: automatic repair of conditional statement bugs in java programs. *IEEE Trans Softw Eng* 43(1):34–55

- Yadav P, Tam D, Choshen L, Raffel CA, Bansal M (2024) Ties-merging: resolving interference when merging models. *Advances in Neural Information Processing Systems* p 36
- Yang E, Wang Z, Shen L, Liu S, Guo G, Wang X, Tao D (2023) Adamerging: adaptive model merging for multi-task learning. arXiv preprint [arXiv:2310.02575](https://arxiv.org/abs/2310.02575)
- Yang E, Shen L, Guo G, Wang X, Cao X, Zhang J, Tao D (2024) Model merging in llms, mllms, and beyond: methods, theories, applications and opportunities. [arXiv:2408.07666](https://arxiv.org/abs/2408.07666) arXiv preprint
- Yu L, Yu B, Yu H, Huang F, Li Y (2023) Language models are super mario: absorbing abilities from homologous models as a free lunch. arXiv preprint [arXiv:2311.03099](https://arxiv.org/abs/2311.03099)
- Yuan C, Hu W, Tian G, Yang S, Wang H (2013) Multi-task sparse learning with beta process prior for action recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 423–429
- Zhang Y, Yang Q (2018) An overview of multi-task learning. *Natl Sci Rev* 5(1):30–43
- Zhang J, Liu J, He J et al (2023) Composing parameter-efficient modules with arithmetic operation. *Adv Neural Inf Process Syst* 36:12589–12610
- Zhao Y, Zhang W, Wang H, Kawaguchi K, Bing L (2024) Adamergex: cross-lingual transfer with large language models via adaptive adapter merging. [arXiv:2402.18913](https://arxiv.org/abs/2402.18913) arXiv preprint
- Zhuo TY, Zebaze A, Suppattarachai N, von Werra L, de Vries H, Liu Q, Muennighoff N (2024) Astraios: parameter-efficient instruction tuning code large language models. arXiv preprint [arXiv:2401.00788](https://arxiv.org/abs/2401.00788)
- Zou W, Shen Z, Li Q, Ge J, Li C, Chen X, Shen X, Huang L, Luo B (2025) Experimental evaluation of parameter-efficient fine-tuning for software engineering tasks. *ACM Transactions on Software Engineering and Methodology*

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Meghdad Dehghan¹ · Jie JW Wu²  · Fatemeh H. Fard³ · Ali Ouni⁴

✉ Jie JW Wu
jie.jw.wu@mtu.edu

Meghdad Dehghan
meghdad.dehghan@ubc.ca

Fatemeh H. Fard
fatemeh.fard@ubc.ca

Ali Ouni
ali.ouni@etsmtl.ca

¹ Department of Computer Science, University of British Columbia, Kelowna, B.C. V1V 1V7, Canada

² Department of Computer Science, Michigan Technological University, Houghton, MI, USA

³ Department of Computer Science, University of British Columbia, Kelowna, B.C. V1V 1V7, Canada

⁴ Department of Software and IT Engineering, École de technologie supérieure (ÉTS), Montreal, QC, Canada