

FusionBench: A Unified Library and Comprehensive Benchmark for Deep Model Fusion

Anke Tang¹

ANKETANG@WHU.EDU.CN

Li Shen^{2*}

MATHSHENLI@GMAIL.COM

Yong Luo^{1*}

YLUO180@GMAIL.COM

Enneng Yang²

ENNENGYANG@GMAIL.COM

Han Hu³

HHU@BIT.EDU.CN

Lefei Zhang¹

ZHANGLEFEI@WHU.EDU.CN

Bo Du^{1*}

DUBO@WHU.EDU.CN

Dacheng Tao⁴

DACHENG.TAO@NTU.EDU.SG

¹*School of Computer Science, National Engineering Research Center for Multimedia Software and Hubei Key Laboratory of Multimedia and Network Communication Engineering, Wuhan University, Wuhan, China* ²*Shenzhen Campus of Sun Yat-sen University, China* ³*Beijing Institute of Technology, Beijing, China* ⁴*Nanyang Technological University, Singapore*

Editor: Zeyi Wen

Abstract

Deep model fusion is an emerging technique that unifies the predictions or parameters of several deep neural networks into a single better-performing model in a cost-effective and data-efficient manner. Although a variety of deep model fusion techniques have been introduced, their evaluations tend to be inconsistent and often inadequate to validate their effectiveness and robustness. We present *FusionBench*, the first benchmark and a unified library designed specifically for deep model fusion. Our benchmark consists of multiple tasks, each with different settings of models and datasets. This variety allows us to compare fusion methods across different scenarios and model scales. Additionally, FusionBench serves as a unified library for easy implementation and testing of new fusion techniques. FusionBench is open source and actively maintained, with community contributions encouraged.

 Code Repository: https://github.com/tanganke/fusion_bench

 Online Documentation: https://tanganke.github.io/fusion_bench

 HuggingFace Collections (models & datasets): <https://huggingface.co/tanganke>

Keywords: Deep Model Fusion, Model Merging, Multi-Task Learning, Benchmark

1. Introduction

In recent years, a new paradigm called “learn from model” has emerged in the field of deep learning, which focuses on leveraging the knowledge embedded in existing models to develop new ones (Zheng et al., 2025). This paradigm has been widely adopted in various scenarios,

*. Corresponding authors: Li Shen (mathshenli@gmail.com), Yong Luo (yluo180@gmail.com) and Bo Du (dubo@whu.edu.cn).

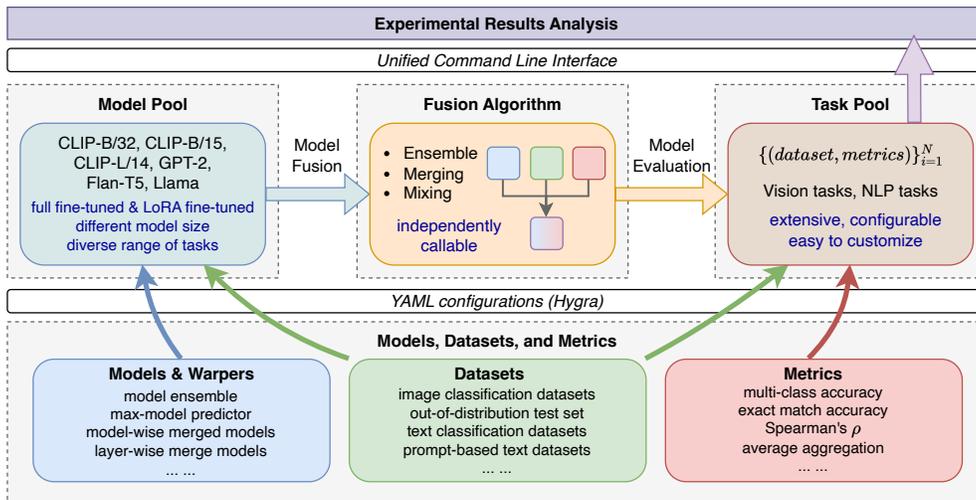


Figure 1: The general framework of the modularized FusionBench codebase.

such as model tuning (He et al., 2022; Chung et al., 2024), model distillation (Hinton, 2015), model pruning (Han et al., 2016; Asif et al., 2020), model editing (Mitchell et al., 2021; Zhang et al., 2024), and so on. Among these methods, deep model fusion is particularly appealing. It merges the parameters or predictions of multiple models to create a more robust and efficient unified model. Due to its effectiveness and scalability, many new techniques for deep model fusion have recently been proposed (Li et al., 2023).

Deep model fusion offers both scalability and data efficiency by utilizing the knowledge embedded in pre-existing models, rather than requiring training from scratch. This approach significantly accelerates model development, making it a practical solution in the current era dominated by large foundation models. Despite its potential, the evaluation of deep model fusion techniques often suffers from inconsistency and inadequacy. Current libraries like MergeKit (Goddard et al., 2024) are primarily focused only on large language model merging and do not provide evaluation frameworks. Standardized assessments across different tasks and models are lacking, making it challenging to evaluate the effectiveness and robustness of deep model fusion techniques. The potential reasons for this inconsistency include the rapid development of new techniques, the absence of standardized tasks and models, and the variety of settings (such as different fine-tuning strategies). Additionally, challenges in implementing or replicating prior work contribute to these inconsistencies. A more detailed comparison with existing model fusion methods and toolkits is provided in Section A.

To tackle these challenges, we develop the first comprehensive benchmark dedicated to deep model fusion, called *FusionBench*. FusionBench is designed as a modular and flexible framework with three primary components: the *Algorithm Module*, the *Model Pool Module*, and the *Task Pool Module* as shown in Figure 1. Each component offers full configurability, enabling straightforward integration of novel algorithms, models, and evaluation tasks. Through Hydra-based configuration management, all experimental setups are streamlined via a single unified interface. Furthermore, FusionBench provides extensive documentation, hands-on tutorials, illustrative code samples, and ready-to-use model collections.

2. FusionBench

FusionBench is available for installation through two simple methods: using pip with ‘`pip install fusion-bench`’ or cloning the repository directly with ‘`git clone`’. The general framework of the modularized FusionBench codebase is shown in Figure 1, which consists of three primary elements: *Algorithm Module*, *Model Pool Module*, and *Task Pool Module*. In Section 2.1, we introduce the three major components, which is designed to be flexible and modular, allowing users to easily run experiments and evaluate the performance of model fusion algorithms. In Section 2.2, we introduce the implemented model fusion algorithms. Finally, in Section 2.3, we discuss the documentation and tutorials provided to help users understand the benchmark and effectively use the codebase. In Appendix C, we provide a flowchart to illustrate the process of running experiments and evaluating the merged models.

Main Program Pipeline. The typical pipeline consists of the following steps: (1) Users configure the experiment by selecting a fusion algorithm, models to merge, and optional evaluation tasks. (2) The system loads the specified models, algorithm, and evaluation tasks. (3) The algorithm merges the selected models, with an option to save the resulting model. (4) When evaluation tasks are specified, the merged model is tested and results are reported.

2.1 Components

As shown in Figure 1, the codebase is composed of three primary components: *Algorithm Module*, *Model Pool Module*, and *Task Pool Module*, which are responsible for implementing the model fusion algorithms, managing the models and datasets used during the algorithm execution, and managing the tasks and evaluation metrics to be evaluated, respectively. Additionally, we provide a unified command line interface (CLI) ‘`fusion_bench`’ to facilitate the use of the codebase and to enable users to easily run experiments.

- **Algorithm Module** is the core component, which contains the implementation of all model fusion algorithms. Each algorithm is implemented as a separate Python class inheriting from the base class `fusion_bench.BaseAlgorithm`. The algorithm classes are flexible - they can be used through our CLI or in custom code.
- **Model Pool Module** provides a unified interface for managing models and datasets.
- **Task Pool Module** evaluates model performance on different tasks, where each task combines a dataset with specific evaluation metrics.

2.2 Implemented Algorithms

Our benchmark incorporates a comprehensive collection of cutting-edge model fusion algorithms, spanning three key categories: ensemble methods, merging methods, and mixing methods. Additionally, we have recently extended FusionBench to include model compression methods, which, while not strictly fusion techniques, are complementary approaches that can be applied to reduce the computational overhead of models or compress individual models before and after fusion. Additionally, we evaluated their prominence in the research literature and their practical utility. The complete listing of these implemented algorithms can be found in Table 1 and FusionBench is under active development and will be updated to include the latest and most effective model fusion algorithms.

Table 1: Implemented algorithms in FusionBench.

Method	Requirement
<i>Model Ensemble Methods</i>	
Simple Ensemble (Sagi and Rokach, 2018)	-
Weighted Ensemble (Sagi and Rokach, 2018)	hyperparameter search
Max-Model Predictor (Wu et al., 2019)	-
<i>Model Merging Methods</i>	
Simple Average / Modelsoups (Wortsman et al., 2022)	-
Weighted Average (Matena and Raffel, 2022)	hyperparameter search
Fisher Merging (Matena and Raffel, 2022)	compute weights on labeled data
RegMean (Jin et al., 2023)	compute weights on labeled data
RegMean++ (Nguyen et al., 2025)	compute weights on labeled data
Concrete Subspace (Tang et al., 2023b)	test-time adaptation training
Task Arithmetic (Ilharco et al., 2023)	hyperparameter search
Ties-Merging (Yadav et al., 2023)	hyperparameter search
Task-Wise AdaMerging (Yang et al., 2024b)	test-time adaptation training
Layer-Wise AdaMerging (Yang et al., 2024b)	test-time adaptation training
Representation Surgery (Yang et al., 2024a)	unlabeled data samples
TALL mask (Wang et al., 2024)	-
Task Singular Vectors (Gargiulo et al., 2025)	-
Isotropic Merging (Marczak et al., 2025)	-
OPCM (Tang et al., 2025b)	-
FW-Merging (Chen et al., 2025)	labeled data samples
RanDeS (Zhou et al., 2025)	-
<i>Model Mixing Methods</i>	
MoE-based Upscaling (Komatsuzaki et al., 2023)	pre-training to recover performance
MoE-based Merging (Komatsuzaki et al., 2023)	training on the combined model
Depth Upscaling (Kim et al., 2024)	pre-training to recover performance
Model Recombination (Hu et al., 2024)	training on the combined model
Weight-Ensemble MoE (Tang et al., 2024b)	test-time adaptation training, vision tasks
WE-MoE V2 (Shen et al., 2025)	test-time adaptation training, vision tasks
Pareto-Driven Merging (Tang et al., 2024a)	training datasets
SMILE Upscaling (Tang et al., 2025a)	-
<i>Model Compression Methods</i>	
BitDelta (Liu et al., 2024)	calibration data
Magnitude Pruning (Frankle et al., 2020)	-
Wanda Pruning (Sun et al., 2024)	calibration data
SparseGPT Pruning (Frantar and Alistarh, 2023)	calibration data
Expert Pruning (Lu et al., 2024a)	calibration data

2.3 Online Documentation and Tutorials

We provide detailed documentation and tutorials on our project website to help users get started. These resources cover the basics of model fusion, how to run experiments, and how to evaluate results. We also include example pre-run results showing how different fusion algorithms perform across various tasks.

3. Conclusion and Discussion

FusionBench offers a robust framework for evaluating deep model fusion algorithms and a unified library for developing new ones. Its scalable and extensible architecture simplifies the creation of deep model fusion techniques. Furthermore, FusionBench provides curated datasets and models to ensure fair comparisons. With comprehensive documentation and tutorials, it is accessible to both beginners and experienced researchers. We encourage the community to use this benchmark to advance the field of deep model fusion.

Acknowledgments and Disclosure of Funding

This work is supported by the National Natural Science Foundation of China (Grant No. 62225113, U23A20318, U2336211, 62576364 and 62276195), the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (Grant No. JYB2025XDXM704), the Shenzhen Basic Research Project (Natural Science Foundation) Basic Research Key Project (NO. JCYJ20241202124430041), the Foundation for Innovative Research Groups of Hubei Province (Grant No. 2024AFA017) and the Science and Technology Major Project of Hubei Province (Grant No. 2024BAB046). Dr. Tao's research is partially supported by NTU RSR and Start Up Grants. The numerical calculations in this paper have been done on the supercomputing system in the Supercomputing Center of Wuhan University.

The appendix is organized into several sections, each providing additional insights and details related to different aspects of the main work.

A	Related Work of Deep Model Fusion	7
A.1	A Taxonomy of Deep Model Fusion Techniques	7
A.2	Related Benchmarks and Toolkits	8
B	Project Structure	8
C	FusionBench Architecture and Implementation Guide	9
C.1	Implementing Customized Algorithm in FusionBench	9
C.2	Customizing Model Pools in FusionBench	12
D	Fine-Tuned Model Collections in FusionBench	13
D.1	CLIP-ViT Models	14
D.2	ResNet-50 Models	15
D.3	GPT-2 Models	16
D.4	Flan-T5 Models	17
E	Experimental Evaluation and Benchmarking Results	18
E.1	Experimental Setup	18
E.2	Computational Cost Analysis	18
E.3	Multi-Task Deep Model Fusion	18
E.4	Generalization and Robustness Evaluation	22
E.5	Scaling to Large-Scale Neural Networks	26
F	Prompt-Based Text-to-Text Generation	27

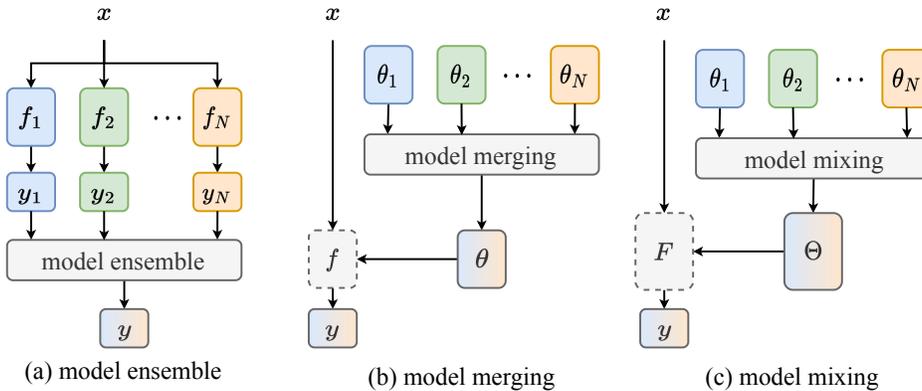


Figure 2: A taxonomy of deep model fusion techniques.

Appendix A. Related Work of Deep Model Fusion

Deep model fusion, often referred to in the literature simply as “model fusion”, aims to combine the outputs or parameters of multiple deep neural networks into a unified model with enhanced performance, efficiency, or robustness. Different researchers may categorize these techniques in various ways based on their understanding and points of view. Here, we propose a taxonomy that divides these techniques into three major categories: *Model Ensemble*, *Model Merging*, and *Model Mixing*. Each of these categories approaches model fusion from a unique perspective, offering distinct advantages and applicability. In the following, we provide detailed explanations, formal definitions, and analyze their strengths and weaknesses. A visualization of the taxonomy is shown in Figure 2.

A.1 A Taxonomy of Deep Model Fusion Techniques

Model Ensemble methods combine the predictions of multiple models to improve the overall performance of a machine learning system (Sagi and Rokach, 2018), where the collective knowledge is often more accurate and reliable than that of any individual model. Mathematically, given a set of N models $\{f_1, f_2, \dots, f_N\}$, which can be homogeneous or heterogeneous, we use their predictions to obtain a global prediction $y = \mathcal{A}_{ensemble}(x; f_1, f_2, \dots, f_N; w)$, where $\mathcal{A}_{ensemble}$ is an ensemble algorithm and w are the algorithmic parameters. Each model f_i can also be associated with a specification to indicate its weight or importance in the ensemble (Pathak et al., 2010; Zhou, 2016; Wu et al., 2021; Tang et al., 2023a). Ensemble methods are widely used and effective in improving performance but are often expensive to use and manage. Recent research has also investigated efficient techniques for model ensembles (Wen et al., 2019; Chen et al., 2024; Allingham et al., 2021).

Model Merging methods integrate the parameters of multiple models into a unified model, enhancing efficiency in terms of inference cost and storage, and enabling scalable model fusion. Given a set of N isomorphic models $\{f_i(\cdot; \theta_i)\}_{i=1}^N$, each parameterized with θ_i , we merge them into a single model with parameters $\theta = \mathcal{A}_{merging}(\theta_1, \theta_2, \dots, \theta_N; w)$, where $\mathcal{A}_{merging} : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}^d$ is a merging algorithm and w are the algorithmic parameters. The merged model can be expressed as $f(\cdot; \theta)$. This method can be implemented through linear interpolation in parameter space (Wortsman et al., 2022; Ilharco et al., 2023; Yadav et al., 2023; Matena and Raffel, 2022; Yu et al., 2024; Chronopoulou et al., 2023; Rame et al., 2024; Ortiz-Jimenez et al., 2024; Liu and Soatto, 2023), leveraging mode connectivity (Draxler et al., 2018; Frankle et al., 2020; Benton et al., 2021; Garipov et al., 2018; Qu and Horvath, 2024), aligning features, parameters or gradients (Liu et al., 2022a; Ainsworth et al., 2023; Jin et al., 2023; Tam et al., 2024; Stoica et al., 2023; Jang et al., 2024; Daheim et al., 2024; Yang et al., 2024a), subspace-based methods (Tang et al., 2023b; Wang et al., 2024; Yi et al., 2024; Zhu et al., 2024; Xu et al., 2024), and ensemble distillation (Wan et al., 2023, 2025). Model merging methods are often performed in a data-efficient manner, the algorithmic parameters w can also be learned during test time via test-time adaptation (TTA) training or meta-learning for a more seamless merging (Yang et al., 2024b; Tang et al., 2023b).

Model Mixing methods fuse the components of multiple models to create a new heterogeneous model, which can be more flexible and adaptive than the original models. Mathematically, given a set of N models $\{f_i(\cdot; \theta_i)\}_{i=1}^N$, each parameterized with $\theta_i \in \mathbb{R}^d$, we mix their components to obtain a new model with parameters $\Theta = \mathcal{A}_{mixing}(\theta_1, \theta_2, \dots, \theta_N; w) \in \mathbb{R}^{d'}$,

where $\mathcal{A}_{mixing} : \mathbb{R}^{N \times d} \mapsto \mathbb{R}^d$ is a mixing algorithm and w is the algorithmic parameters. The mixed model can be expressed as $F(\cdot; \Theta)$, which often has more parameters than the original models, and thus can be more expressive and powerful to capture the underlying patterns in the data. Model mixing methods can be implemented through layer recombinations (Hu et al., 2024; Jiang, 2024), model stitching (Lenc and Vedaldi, 2015; Moschella et al., 2023), or upscale to create a Mixture of Experts (MoE)-based sparse model (Komatsuzaki et al., 2023; Ye and Xu, 2023; Tang et al., 2024b; Lu et al., 2024b; Dai et al., 2024; Zhao et al., 2024; Ostapenko et al., 2024; Tang et al., 2025a; Yadav et al., 2024).

A.2 Related Benchmarks and Toolkits

Although several model fusion methods have been proposed, comprehensive benchmarks and unified toolkits are still lacking in this field. A recent notable work, MergeKit (Goddard et al., 2024), provides a collection of model fusion techniques *specifically designed for merging large language models (LLMs)*, focusing primarily on methods for Transformer-based architectures. While MergeKit excels in providing practical tools for LLM merging, its scope is inherently limited to a single domain (natural language processing) and a specific model family (Transformers). Similarly, MergeBench (He et al., 2025) is targeted exclusively at evaluating LLM merging techniques, providing benchmarks only within the language modeling domain.

In contrast, FusionBench addresses several critical gaps in the current landscape, making FusionBench a more generalized and versatile platform for systematically assessing the performance of model fusion approaches across various domains and architectures: (1) *Broader Domain Coverage*: Unlike existing tools that focus solely on NLP, FusionBench spans both computer vision and natural language processing domains, enabling cross-domain evaluation and comparison. (2) *Comprehensive Taxonomy*: While MergeKit and MergeBench concentrate primarily on model merging methods, FusionBench encompasses a complete taxonomy including model ensemble, model merging, and model mixing methods, providing a unified framework for all major fusion paradigms. (3) *Research-Oriented Evaluation*: FusionBench is designed as a research platform that includes diverse fine-tuned models, standardized evaluation protocols, and comparative analysis utilities across different architectures (CNNs, Vision Transformers, CLIP models, etc.) and tasks. (4) *Algorithmic Diversity*: Beyond basic merging techniques, FusionBench incorporates advanced methods such as subspace-based fusion, gradient alignment, and mixture-of-experts approaches.

Appendix B. Project Structure

To provide a clear overview of FusionBench’s organization, we present its comprehensive project structure in Listing 1. This structure reflects the modular design philosophy of our benchmark, with distinct directories for configuration management, documentation, example implementations, core functionality, and testing.

```

1 |-- config/                # Configuration files managed by Hydra
2 |   |-- method/           # Configs for implemented algorithms
3 |   |-- modelpool/        # Configs for model pool
4 |   |-- taskpool/         # Configs for task pool

```

```

5 | |-- model/           # Configs for models
6 | |-- dataset/       # Configs for datasets
7 | |-- docs/          # Documentation (mkdocs)
8 | |-- examples/      # Example scripts and notebooks
9 | |-- {method_name}/ # Method-specific examples
10 | |-- fusion_bench/  # Main package
11 | |-- method/        # Fusion algorithm implementations
12 | | |-- {method_name}/ # Method-specific code
13 | |-- modelpool/     # Model pool implementations
14 | |-- taskpool/     # Task pool implementations
15 | |-- ...           # Models, datasets, and other utilities
16 | |-- tests/        # Unit tests

```

Listing 1: Project structure of FusionBench.

Appendix C. FusionBench Architecture and Implementation Guide

Figure 3 illustrates the steps to fuse a model in FusionBench across three main stages: creating a model pool, instantiating and running an algorithm, and (optional) model evaluation.

Pseudocode Example. The pseudocode for the pipeline is shown in Algorithm 1.

Algorithm 1 Pseudocode for FusionBench main pipeline

- 1: Load and parse configuration from YAML files
 - 2: Instantiate *fusion_algorithm*, *model_pool*, and (optionally) *task_pool*
 - 3: $merged_model \leftarrow fusion_algorithm.run(model_pool)$
 - 4: Save *merged_model* if save path is specified
 - 5: **if** task pool is specified **then**
 - 6: $report \leftarrow task_pool.evaluate(merged_model)$
 - 7: Save *report* if save path is specified
 - 8: **end if**
-

This design allows users to flexibly experiment with different algorithms, models, and tasks by simply modifying the configuration, without changing the core program logic.

C.1 Implementing Customized Algorithm in FusionBench

Here, we provide an example of implementing a customized algorithm in Listing 2. The derived algorithm class should inherit from the “BaseAlgorithm” class and implement the “run” method. The “run” method takes a “BaseModelPool” object as input, which is responsible for managing the models and dataset to be loaded. The “run” method should be implemented to fuse the models and return the fused model.

```

1 from fusion_bench import BaseAlgorithm, BaseModelPool
2
3 class DerivedAlgorithm(BaseAlgorithm):
4     """
5     An example of a derived algorithm class.

```

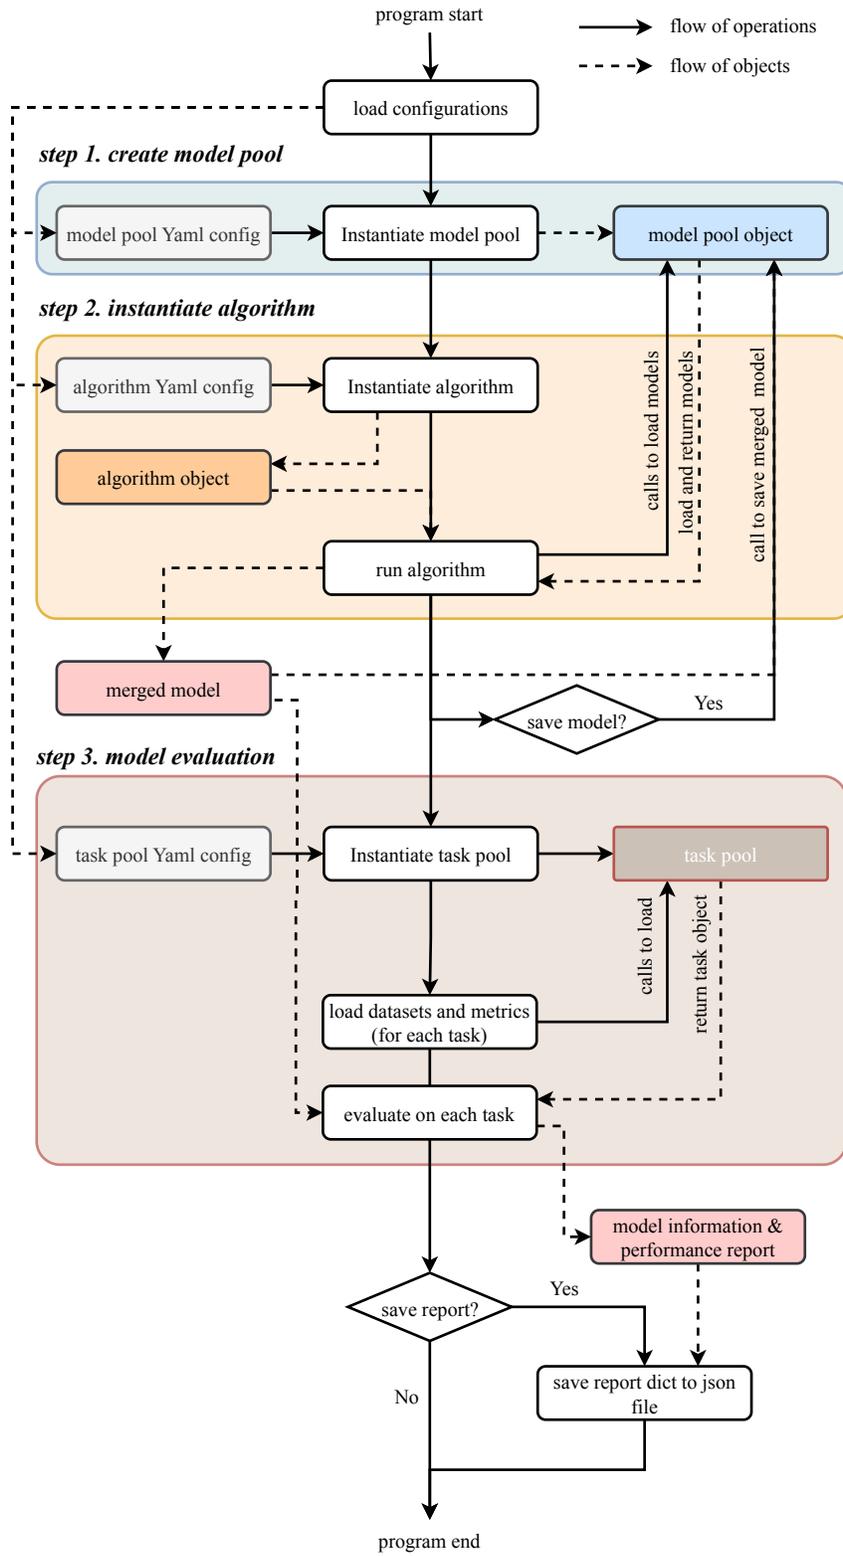


Figure 3: Flowchart of FusionBench.

```

6      """
7
8      # _config_mapping maps the attribution to the corresponding key in
9      # the configuration file.
10     # this is optional and can be used to serialize the object to a
11     # configuration file.
12     # 'self.config.hyperparam_1' will be mapped to the attribute
13     # 'hyperparam_attr_1'.
14     _config_mapping = BaseModelFusionAlgorithm._config_mapping | {
15         "hyperparam_attr_1": "hyperparam_1",
16         "hyperparam_attr_2": "hyperparam_2",
17     }
18
19     def __init__(self, hyperparam_1, hyperparam_2, **kwargs):
20         super().__init__(**kwargs)
21         self.hyperparam_attr_1 = hyperparam_1
22         self.hyperparam_attr_2 = hyperparam_2
23
24     def run(self, modelpool: BaseModelPool):
25         # modelpool is an object that responsible for managing the models
26         # and dataset to be loaded.
27         # implement the fusion algorithm here.
28         raise NotImplementedError(
29             "DerivedAlgorithm.run() is not implemented."
30         )

```

Listing 2: Implementing new algorithm in FusionBench.

As an illustration, we can implement a deep model fusion algorithm that merges models through parameter averaging (Wortsman et al., 2022). The following code snippet demonstrates this approach (Listing 3):

```

1  class SimpleAverageAlgorithm(
2      BaseAlgorithm,
3      SimpleProfilerMixin,
4  ):
5      @torch.no_grad()
6      def run(self, modelpool: Union[BaseModelPool, Dict[str, nn.Module]]):
7          """
8          Fuse the models in the given model pool using simple averaging.
9
10         This method iterates over the names of the models in the model
11         pool, loads each model, and appends it to a list.
12         It then returns the simple average of the models in the list.
13
14         Args:
15             modelpool: The pool of models to fuse.
16
17         Returns:
18             The fused model obtained by simple averaging.
19         """

```

```

19     if isinstance(modelpool, dict):
20         modelpool = BaseModelPool(modelpool)
21
22     log.info(
23         f"Fusing models using simple average on
24         {len(modelpool.model_names)} models."
25         f"models: {modelpool.model_names}"
26     )
27     sd: Optional[StateDictType] = None
28     forward_model = None
29     merged_model_names = []
30
31     for model_name in modelpool.model_names:
32         with self.profile("load model"):
33             model = modelpool.load_model(model_name)
34             merged_model_names.append(model_name)
35             print(f"load model of type: {type(model).__name__}")
36             with self.profile("merge weights"):
37                 if sd is None:
38                     # Initialize the state dictionary with the first
39                     # model's state dictionary
40                     sd = model.state_dict(keep_vars=True)
41                     forward_model = model
42                 else:
43                     # Add the current model's state dictionary to the
44                     # accumulated state dictionary
45                     sd = state_dict_add(sd,
46                     model.state_dict(keep_vars=True))
47             with self.profile("merge weights"):
48                 # Divide the accumulated state dictionary by the number of
49                 # models to get the average
50                 sd = state_dict_div(sd, len(modelpool.model_names))
51
52     forward_model.load_state_dict(sd)
53     # print profile report and log the merged models
54     self.print_profile_summary()
55     log.info(f"merged {len(merged_model_names)} models:")
56     for model_name in merged_model_names:
57         log.info(f"  - {model_name}")
58     return forward_model

```

Listing 3: Implementation of simple average algorithm in FusionBench.

C.2 Customizing Model Pools in FusionBench

The *Model Pool Module* serves as an abstraction layer that manages model loading, saving, and dataset handling during fusion operations. To extend FusionBench to new model architectures, users need to implement a custom model pool class that inherits from `BaseModelPool` and defines architecture-specific operations. This design principle ensures that fusion algorithms

remain architecture-agnostic and can be applied universally across different model types without modification.

Extensibility to Modern Architectures. FusionBench’s modular architecture facilitates straightforward integration of emerging model architectures. The key advantage of this design is that once a model pool class is implemented for a particular architecture, all existing fusion algorithms in the benchmark become immediately applicable to that architecture. This extensibility has enabled us to expand beyond our initial focus on established models (CLIP-ViT, ResNet-50, GPT-2, and Flan-T5) to encompass state-of-the-art architectures across multiple domains.

For computer vision tasks, we have integrated modern architectures including ConvNeXt (Liu et al., 2022b) and DINOv2 (Oquab et al., 2024), both of which represent significant advances in efficient visual representation learning. These models can be loaded through dedicated model pool classes (`ConvNeXtForImageClassificationPool` and `DINOv2ForImageClassificationPool`) with comprehensive online documentation available on our project website.

For natural language processing tasks, our `CausalLMPool` class provides broad compatibility with contemporary large language model architectures. The `CausalLMPool` leverages the HuggingFace Transformers library, allowing seamless integration of new architectures as they are released.

Appendix D. Fine-Tuned Model Collections in FusionBench

This section outlines the comprehensive model collections provided in FusionBench, including their experimental configurations and performance evaluation results across various single-task specializations.

1. **CLIP-ViT Models:** In FusionBench, we provide comprehensive CLIP-ViT model collections fine-tuned on up to 20 image classification tasks. The most commonly used eight tasks are SUN397 (Xiao et al., 2010), Cars (Krause et al., 2013), RESISC45 (Cheng et al., 2017), EuroSAT (Helber et al., 2018), SVHN (Netzer et al., 2011), GTSRB (Stallkamp et al., 2012), MNIST (Lecun et al., 1998), and DTD (Cimpoi et al., 2014). Additionally, we extend our collection to include 12 more tasks: Oxford Flowers102 (Nilsback and Zisserman, 2008), PatchCamelyon (Veeling et al., 2018), FER2013 (Goodfellow et al., 2013), Oxford-IIIT Pet (Parkhi et al., 2012), STL-10 (Coates et al., 2011), CIFAR-100 (Krizhevsky et al., 2009), CIFAR-10, Food-101 (Bossard et al., 2014), Fashion-MNIST (Xiao et al., 2017), EMNIST Letters (Cohen et al., 2017), Kuzushiji-MNIST (Clanuwat et al., 2018), and Rendered-SST2 (Socher et al., 2013), covering diverse domains from natural images to histopathologic data and facial expressions. We use the Adam Optimizer with a fixed learning rate of $1e-5$ over 4000 training steps (batch size=32). Only the vision encoder is fine-tuned while the text encoder remains fixed to preserve the open-vocabulary property of the model. All fine-tuned models are publicly available on the HuggingFace Model Hub ¹ and Modelscope ² for easy access and reproducibility. In Section D.1, we present the

1. <https://huggingface.co/>

2. <https://www.modelscope.cn/home>

- performance of fine-tuned single-task CLIP-ViT-B/32 and CLIP-ViT-L/14 models on the eight most commonly used eight image classification tasks.
2. **ResNet-50 Models:** We fine-tune ResNet-50 models on three scene understanding tasks: segmentation, depth estimation, and normal estimation using the NYUv2 dataset, each with a learning rate of $1e-4$ for 40 epochs, the learning rate is reduced by a factor of 0.5 every 10 epochs. The performance of fine-tuned single-task ResNet-50 models on the NYUv2 dataset is shown in Section D.2.
 3. **GPT-2 Models:** We fine-tune GPT-2 models on seven text classification tasks from the GLUE benchmark (Wang et al., 2018) to evaluate their effectiveness in natural language understanding scenarios. The fine-tuning process employs the Adam optimizer with a constant learning rate of $5e-5$ over 3 epochs, following standard practices for transformer-based language models. These tasks include CoLA, MNLI, MRPC, QNLI, QQP, RTE, and SST-2. The performance evaluation results of these fine-tuned single-task GPT-2 models are presented in Section D.3.
 4. **Flan-T5 Models:** We fine-tune both Flan-T5-base and Flan-T5-large models on eight text-to-text generation tasks from the GLUE benchmark. For reproducibility and transparency, we provide the detailed prompt templates used for each task in Section F, ensuring that researchers can replicate our experimental setup and build upon our work. The performance evaluation results of these LoRA fine-tuned Flan-T5-base and Flan-T5-large models across all eight tasks are presented in Section D.4.

Based on the performance metrics detailed in Section D.1, D.2, D.3 and D.4, we observe that the fine-tuned models demonstrate high accuracy on specific tasks. This observation holds true across various model architectures and task domains, indicating the effectiveness of the fine-tuning process in adapting pre-trained models to excel in particular applications. Furthermore, fine-tuning a model on one task can lead to both positive or negative transfer effects on other tasks. Positive transfer occurs when the knowledge gained from the target task enhances the model’s performance on another task, while negative transfer arises when the fine-tuning process on the target task hinders the model’s ability to perform on others.

D.1 CLIP-ViT Models

The performance of fine-tuned CLIP-ViT-B/32 and CLIP-ViT-L/14 models on the eight most commonly used image classification tasks is shown in Tables 2 and 3, respectively. In Figure 4, we visualize the cosine similarity matrices of task vectors for CLIP-ViT-B/32 and CLIP-ViT-L/14 models. Note that while we present results for the eight core tasks in this paper for conciseness, comprehensive evaluation results for all 20 tasks across CLIP-ViT-B/32, CLIP-ViT-B/16, and CLIP-ViT-L/14 models are available in our online documentation. Following (Ilharco et al., 2023), the task vector is defined as the difference between a fine-tuned model’s parameters and the original pre-trained model’s parameters: $\tau = \theta_{ft} - \theta_{pre}$, where θ_{ft} represents the fine-tuned weights and θ_{pre} represents the pre-trained weights. We note that the task vectors for models from various tasks are nearly orthogonal. This suggests that the knowledge specific to each task resides in distinct directions or subspaces. This finding motivates the exploration of locating subspaces in which the knowledge of different tasks can be merged effectively, as discussed in Tang et al. (2023b); Wang et al. (2024).

Table 2: Performance of fine-tuned single-task CLIP-ViT-B/32 models on the eight image classification tasks.

Model	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD
Pre-trained	63.2	59.8	60.7	46.0	31.6	32.5	48.2	43.9
SUN397	75.0	47.0	54.3	46.5	28.3	26.4	44.3	41.6
Cars	56.6	78.3	50.9	38.4	30.2	30.6	49.7	41.8
RESISC45	52.0	47.2	95.2	56.9	23.9	24.3	39.7	35.9
EuroSAT	49.0	39.9	33.5	99.0	11.8	22.9	33.8	35.5
SVHN	40.5	36.3	18.9	9.8	97.3	27.3	81.8	23.2
GTSRB	36.9	33.0	20.6	21.3	41.2	98.9	30.9	23.9
MNIST	50.3	40.0	31.3	17.7	50.1	19.3	99.6	30.7
DTD	54.6	51.3	36.8	25.0	28.9	21.8	47.3	79.7

Table 3: Performance of fine-tuned single-task CLIP-ViT-L/14 models on the eight image classification tasks.

Model	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD
Pre-trained	68.3	77.8	71.0	58.9	58.4	50.6	76.4	55.5
SUN397	82.8	68.4	58.1	49.9	55.0	46.3	79.5	52.8
Cars	67.8	92.9	68.7	56.4	51.7	47.7	80.5	55.6
RESISC45	65.6	69.0	97.4	64.3	38.3	46.6	77.7	49.9
EuroSAT	65.2	69.0	40.6	99.2	33.4	45.6	73.5	47.1
SVHN	66.5	69.0	54.0	19.7	97.9	48.7	92.2	50.1
GTSRB	63.4	64.8	38.7	19.6	71.0	99.2	75.1	45.8
MNIST	56.1	49.8	53.5	26.6	48.2	33.1	99.8	47.1
DTD	66.8	75.3	65.5	43.7	49.5	45.0	68.5	85.5

D.2 ResNet-50 Models

Table 4: Single-task performance of fine-tuned ResNet-50 models on NYUv2 dataset.

Method	SEGMENTATION		DEPTH ESTIMATION		NORMAL
	mIoU \uparrow	Pix Acc \uparrow	Abs Err \downarrow	Rel Err \downarrow	Mean \downarrow
Segmentation	52.0	73.8	242.8	88.7	82.8
Depth Estimation	2.3	6.2	42.5	17.7	82.8
Normal	2.0	4.9	264.0	98.1	24.7

We fine-tune ResNet-50 models on three distinct scene understanding tasks using the NYUv2 dataset: semantic segmentation, depth estimation, and surface normal estimation. Each model is trained with a learning rate of $1e-4$ for 40 epochs, with the learning rate reduced by a factor of 0.5 every 10 epochs. Table 4 presents the single-task performance of

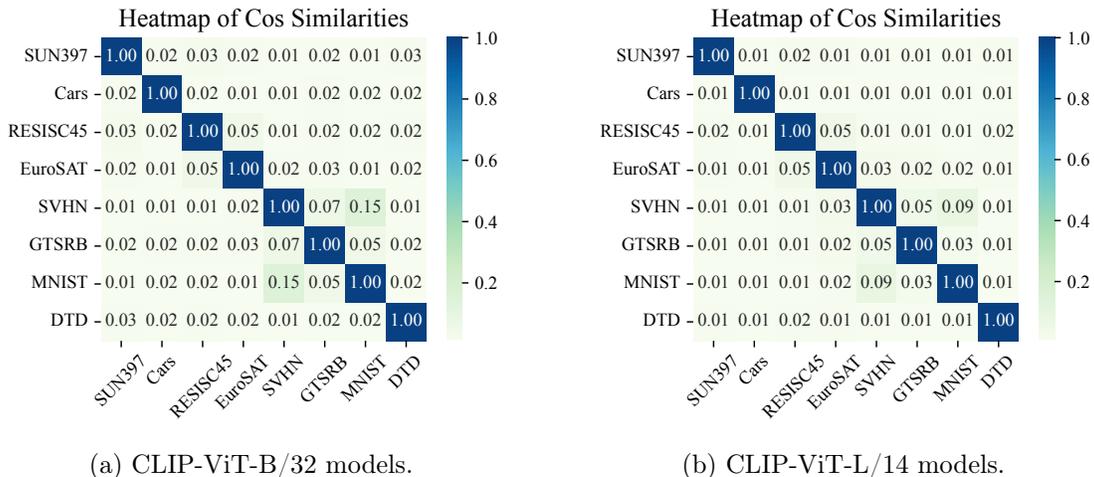


Figure 4: Cosine similarity of task vectors for CLIP-ViT-B/32 and CLIP-ViT-L/14 models.

these specialized models, where each row represents a model trained exclusively on one task and evaluated across all three tasks using task-specific metrics. The results demonstrate clear task specialization: the segmentation model achieves optimal performance on segmentation metrics (52.0 mIoU, 73.8 Pixel Accuracy), the depth estimation model excels at depth-related metrics (42.5 Absolute Error, 17.7 Relative Error), and the normal estimation model performs best on surface normal prediction (24.7 Mean angular error). Notably, models trained on one task show severely poor performance on other tasks, highlighting the challenge of multi-task learning in computer vision applications.

D.3 GPT-2 Models

Table 5: Performance of fine-tuned GPT-2 models on the seven text classification tasks.

Model	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST-2
CoLA	76.8	32.8	68.4	50.4	39.2	48.0	51.0
MNLI	59.5	82.1	33.8	46.5	24.9	57.4	40.5
MRPC	30.8	25.9	80.4	47.1	65.9	49.1	49.1
QNLI	58.7	38.9	30.6	88.3	39.9	48.7	47.0
QQP	31.4	25.7	62.3	45.0	89.6	49.1	49.1
RTE	52.8	47.7	37.5	53.5	33.7	65.3	54.9
SST-2	51.8	32.9	40.2	49.8	56.8	44.4	91.2

We evaluate GPT-2 models fine-tuned on seven distinct text classification tasks from the GLUE benchmark, as detailed in the enumeration above. Table 5 presents the cross-task performance evaluation, where each row represents a model fine-tuned on a specific task (indicated by the row label) and columns show its performance across all seven tasks. The results demonstrate clear task specialization, with each fine-tuned model achieving its highest performance on its target task (diagonal elements).

D.4 Flan-T5 Models

Table 6: Performance of LoRA fine-tuned Flan-T5-Base models on the eight text-to-text generation tasks from the GLUE benchmark.

Model	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST2	STSB
Pre-trained	69.1	56.5	76.2	88.4	82.1	80.1	91.2	62.2
CoLA	69.1	39.9	75.2	89.1	81.1	81.9	90.7	54.0
MNLI	69.4	82.7	73.8	89.3	82.0	79.4	90.9	68.1
MRPC	64.0	44.9	85.5	82.6	81.0	69.0	88.6	73.6
QNLI	68.9	52.7	76.7	90.9	82.8	79.8	91.5	68.9
QQP	65.0	54.6	75.7	89.0	84.0	81.6	90.7	75.3
RTE	64.9	51.8	69.4	89.2	79.8	84.5	90.6	70.1
SST2	68.3	56.6	76.0	88.5	83.4	79.8	92.9	62.6
STSB	65.7	1.7	67.4	89.3	80.1	79.8	90.8	87.4

Table 7: Performance of LoRA fine-tuned Flan-T5-Large models on the eight text-to-text generation tasks from the GLUE benchmark.

Model	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST2	STSB
Pre-trained	73.7	56.6	82.4	91.1	85.5	85.6	94.3	87.5
CoLA	80.2	53.9	81.4	90.8	84.5	84.1	93.9	87.1
MNLI	73.7	88.5	77.9	92.4	85.2	87.7	94.4	86.7
MRPC	75.6	52.6	89.2	92.6	84.4	86.3	94.3	86.3
QNLI	73.5	54.5	82.8	94.4	85.8	85.2	93.7	87.1
QQP	74.0	53.8	82.8	92.5	87.2	85.6	94.5	88.3
RTE	75.6	57.5	69.9	92.8	83.8	91.7	94.6	86.0
SST2	73.6	55.3	82.1	91.6	85.5	85.2	95.2	86.9
STSB	73.4	39.3	82.1	92.6	86.1	83.4	94.0	90.9

We evaluate Flan-T5-base and Flan-T5-large models fine-tuned using LoRA on eight text-to-text generation tasks from the GLUE benchmark (Wang et al., 2018). Tables 6 and 7 present the cross-task performance evaluation for both model sizes, where each row represents a model fine-tuned on a specific task and columns show performance across all eight tasks. The results demonstrate that each specialized model achieves its best performance on its target task (diagonal elements), with the Flan-T5-large models generally outperforming their base counterparts across all tasks. Notably, when evaluating fine-tuned models on non-target tasks, we observe both positive and negative transfer effects. For instance, in Flan-T5-base models, the MNLI-specialized model shows positive transfer to QNLI (89.3% vs 88.% pre-trained baseline), while the STSB-specialized model exhibits negative transfer to MNLI (1.7% vs 56.5% pre-trained baseline).

Appendix E. Experimental Evaluation and Benchmarking Results

In this section, we evaluate the performance of multi-task deep model fusion algorithms on a variety of tasks, as well as analyze the generalization and robustness of these algorithms. We also provide an ablation study to investigate the impact of hyperparameter selection. The majority of experiments presented in this section can be executed using a single NVIDIA RTX 3090 GPU with 24GB memory, ensuring cost-effective reproducibility.

E.1 Experimental Setup

In this section, we conduct a series of multi-task deep model fusion experiments on image classification tasks, text classification tasks, and text-to-text generation tasks to evaluate the performance of multi-task deep model fusion algorithms. These tasks are chosen to cover a wide range of NLP and CV tasks, as described in Section E.3.

E.2 Computational Cost Analysis

While functional requirements listed in Table 1 differ across methods, computational resource implications also vary significantly. For instance, ensemble methods incur high inference costs since multiple models must be stored and executed simultaneously. Model merging methods are generally cost-efficient during inference, as the merged model retains the footprint of a single model, though some approaches require additional data-dependent computations during merging. Mixing methods typically expand the model size and necessitate partial re-training, leading to substantial computational overhead.

Practical Implications: For resource-constrained environments, parameter-free methods like parameter average and task arithmetic offer the best cost-performance trade-off. When additional data is available, alignment-based methods such as RegMean (Jin et al., 2023) and RegMean++ (Nguyen et al., 2025) provide strong performance with moderate computational overhead. For applications requiring maximum performance and having sufficient computational resources or access to the test data distribution, test-time adaptive methods like AdaMerging (Yang et al., 2024b) or model mixing approaches like WE-MoE (Tang et al., 2024b) are recommended despite their higher computational costs. Besides, the `LazyStateDict` utility in FusionBench significantly reduces memory requirements for all methods by enabling on-demand parameter loading, making large-scale model fusion practical.

E.3 Multi-Task Deep Model Fusion

We start our evaluation by comparing multi-task model fusion algorithms across multiple experimental settings, including both computer vision and natural language processing tasks:

1. **Open-Vocabulary Image Classification:** Using the CLIP-ViT models fine-tuned on the eight most commonly used tasks. Results for CLIP-ViT-B/32 and CLIP-ViT-L/14 models are presented in Section E.3.1.
2. **Scene Understanding:** Using ResNet-50 models on the NYUv2 dataset for segmentation, depth estimation, and normal estimation tasks. Detailed results are provided in Section E.3.2.
3. **Text Classification:** Detailed results for GPT-2 models on seven text classification tasks are presented in Section E.3.3.

4. **Text-to-Text Generation:** For Flan-T5-base and Flan-T5-large models fine-tuned with LoRA, the results are presented in Section E.3.4.

In Section E.3.1, E.3.2, E.3.3 and E.3.4, we compare the performance of different multi-task model fusion algorithms across various tasks. Pre-trained models’ performance, fine-tuned models’ performance, and traditional multi-task learning (MTL) methods are provided for reference. In Appendix D, we provide a detailed description of these fine-tuned single-task models.

We have the following key observations: (1) Multi-task model fusion usually outperforms pre-trained models, showing it can transfer knowledge from multiple single-task models to enhance performance. Pre-trained models lack task-specific knowledge as they are not fine-tuned for downstream tasks. (2) Adaptive method (AdaMerging) and MoE-based method perform best in multi-task model fusion, showing the effectiveness of adaptive merging and mixture-of-experts approaches. (3) The performance gap between multi-task model fusion and single-task fine-tuned models (STL) is larger for CLIP-ViT-B/32 compared to CLIP-ViT-L/14. This suggests that multi-task model fusion may be more beneficial for smaller models, as they have more room for improvement through knowledge transfer. (4) Traditional MTL outperforms most multi-task model fusion methods, which indicates that traditional MTL is still a strong baseline for multi-task learning, and there is room for improvement in multi-task model fusion algorithms.

E.3.1 OPEN-VOCABULARY IMAGE CLASSIFICATION USING CLIP-ViT MODELS

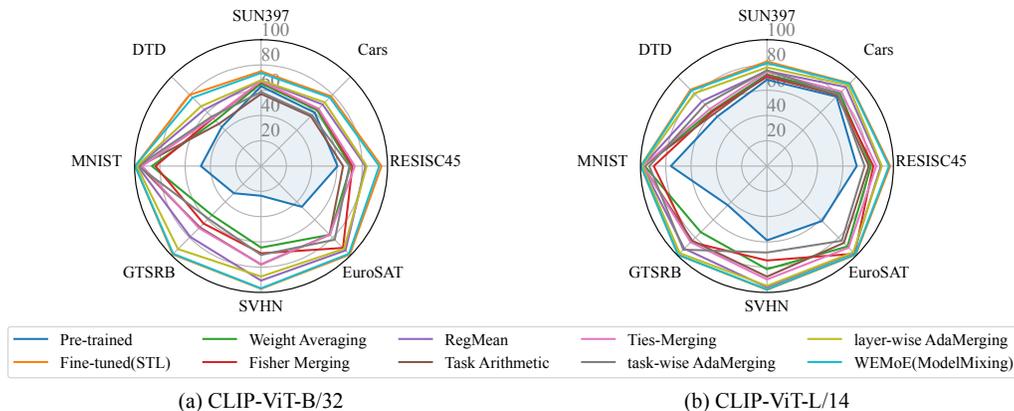


Figure 5: Radar charts comparing the performance of different deep model fusion methods across eight tasks using CLIP-ViT-B/32 and CLIP-ViT-L/14.

We evaluate model fusion methods on open-vocabulary image classification using CLIP-ViT models fine-tuned on eight diverse image classification tasks spanning natural scenes, satellite imagery, handwritten digits, and texture recognition. Tables 8 and 9 present comprehensive results for CLIP-ViT-B/32 and CLIP-ViT-L/14 models respectively, while Figure 5 provides radar chart visualizations highlighting the performance patterns across tasks. The results demonstrate several key insights: (1) Model mixing methods (such as WEMoE and SMILE) consistently achieve the best performance, approaching traditional multi-task learning results with averages of 89.2-89.3% for ViT-B/32 and 93.6-93.8% for

Table 8: Multi-task performance when merging CLIP-ViT-B/32 models on all eight tasks. Here underlined method names indicate methods that require test-time adaptation training, while *italicized method names* indicate model mixing approaches.

Method	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	AVG.
	<i>Reference Methods</i>								
Pre-trained	63.2	59.8	60.7	46.0	31.6	32.5	48.2	43.9	48.2
Individual Fine-tuned	75.0	78.3	95.2	99.0	97.3	98.9	99.6	79.7	90.3
Traditional MTL	72.3	76.6	92.2	97.9	95.5	97.7	99.3	77.7	88.6
	<i>Multi-Task Model Fusion Methods</i>								
Weight Averaging	65.4	62.6	70.8	76.9	64.5	54.9	86.3	50.9	66.5
Fisher Merging	66.7	64.0	72.2	91.6	69.0	64.3	83.5	53.7	70.6
RegMean	68.6	70.0	84.6	95.4	92.6	83.4	98.4	66.1	82.4
RegMean++	69.3	70.5	86.7	96.1	94.1	90.4	99.0	68.7	84.4
Task Arithmetic	57.1	55.7	64.9	76.7	77.9	68.5	96.1	47.2	68.0
Ties-Merging	67.1	64.2	74.1	76.8	77.7	69.4	94.1	54.0	72.2
<u>task-wise AdaMerging</u>	58.6	56.9	69.8	82.4	70.3	58.9	97.2	55.3	68.7
<u>layer-wise AdaMerging</u>	67.9	71.3	83.5	92.7	87.4	92.9	98.2	67.0	82.6
<u>WEMoE (Model Mixing)</u>	73.7	76.8	93.4	98.2	96.8	98.2	99.6	76.6	89.2
<u>SMILE (Model Mixing)</u>	73.6	77.8	92.0	98.3	96.9	98.1	99.6	78.1	89.3

ViT-L/14. (2) Layer-wise adaptive merging (AdaMerging) significantly outperforms task-wise variants, highlighting the importance of fine-grained parameter-level fusion strategies. (3) Simple parameter averaging and Fisher merging provide strong baselines but fall short of more sophisticated approaches like RegMean, which leverages data alignment techniques. (4) The performance gap between fusion methods and individual fine-tuned models is more pronounced for the smaller ViT-B/32 (90.3% vs 89.3%) compared to ViT-L/14 (94.3% vs 93.8%), suggesting that larger models may be more amenable to effective knowledge combination through fusion techniques. This phenomenon has also been discussed in the literature on large language model merging (Yadav et al., 2025).

E.3.2 SCENE UNDERSTANDING USING RESNET-50 MODELS

Figure 6 illustrates the multi-task model fusion problem in scene understanding, where three specialized ResNet-50 models trained individually on segmentation, depth estimation, and normal estimation tasks are merged at the backbone level while preserving task-specific heads. Table 10 presents the results. The results show that single-task ResNet50 models excel only at their own objectives but fail badly on other tasks, highlighting strong specialization without transferability. In contrast, model fusion produces balanced performance across tasks: for example, while fusion methods sacrifice some segmentation accuracy compared to the segmentation expert, they dramatically improve depth and normal performance relative to non-specialists.

Table 10: Experimental results of merging single-task Resnet50 models on three NYUv2 tasks.

Method	SEGMENTATION		DEPTH ESTIMATION		NORMAL
	mIoU \uparrow	Pix Acc \uparrow	Abs Err \downarrow	Rel Err \downarrow	Mean \downarrow
<i>Single-Task Learning</i>					
Segmentation	52.0	73.8	242.8	88.7	82.8
Depth Estimation	2.3	6.2	42.5	17.7	82.8
Normal	2.0	4.9	264.0	98.1	24.7
<i>Multi-Task Model Fusion Methods</i>					
Weight Averaging	39.0	67.0	55.1	22.7	30.4
Task Arithmetic ($\lambda = 0.3$)	33.6	63.3	56.3	23.2	31.3
Ties-Merging ($\lambda = 0.3$)	36.3	61.7	60.5	24.5	33.1

Table 11: Multi-task performance of GPT-2 models on seven text classification tasks.

Method	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST-2	Avg.
<i>Reference Methods</i>								
Fine-tuned (STL)	76.8	82.1	80.4	88.3	89.6	65.3	91.2	82.0
<i>Multi-Task Deep Model Fusion Method</i>								
Simple Average	55.0	55.1	51.0	57.6	76.7	44.8	52.5	56.1
Fisher Merging	54.8	58.0	39.5	63.3	81.5	49.1	64.7	58.7
RegMean	61.7	70.4	65.4	69.7	78.8	56.0	79.7	68.8
Task Arithmetic	68.7	68.6	69.6	70.5	81.8	47.3	83.6	70.0
Ties-Merging	68.4	71.4	68.4	69.6	82.4	47.7	81.8	70.0

and Flan-T5-large models respectively, demonstrating how model fusion performs in the parameter-efficient fine-tuning paradigm. The LoRA adapters are merged into the base models before fusion, creating full fine-tuned models that are then subject to various fusion techniques. Notably, while Flan-T5-base fusion methods show modest improvements over simple averaging (78.2% vs. individual fine-tuned performance of 84.6%), the SMILE model mixing method achieves remarkable performance (84.0%), nearly matching the individual fine-tuned baseline. For Flan-T5-large, all fusion methods demonstrate stronger performance, with the best methods (Task Arithmetic, Ties-Merging, Layer-wise AdaMerging) achieving 87.3-87.6% compared to 89.6% for individual models, suggesting that larger models provide more robust representations for effective knowledge combination across tasks. This trend aligns with observations from the vision experiments, reinforcing the hypothesis that model scale facilitates more effective parameter-space fusion.

E.4 Generalization and Robustness Evaluation

To further assess the generalization and robustness of multi-task model fusion algorithms, we conduct experiments on *unseen tasks* and *corrupted test sets* (or *out-of-distribution test sets*).

Table 12: Experimental results of merging Flan-T5-base (LoRA fine-tuned) models on all eight tasks. Here underlined method names indicate methods that require test-time adaptation training, while *italicized method names* indicate model mixing approaches.

Method	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST2	STSB	Avg.
<i>Reference Methods</i>									
Pre-trained	69.1	56.5	76.2	88.4	82.1	80.1	91.2	62.2	75.7
Individual	69.1	82.7	85.5	90.9	84.0	84.4	92.9	87.4	84.6
<i>Multi-Task Model Fusion Methods</i>									
Weight Averaging	69.7	59.7	78.9	90.1	83.8	80.5	91.2	72.0	78.2
Task Arithmetic	68.8	55.2	78.7	89.8	83.7	79.1	91.5	72.4	77.4
Ties-Merging	68.3	56.3	79.4	89.8	83.7	79.4	91.6	71.2	77.5
Layer-wise AdaMerging	69.1	60.3	78.4	90.0	83.6	79.1	91.6	74.1	78.3
<u>SMILE (Model Mixing)</u>	69.3	82.9	83.8	90.6	83.9	83.4	93.1	85.1	84.0

Table 13: Experimental results of merging Flan-T5-large (LoRA fine-tuned) models on all eight tasks.

Method	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST2	STSB	Avg.
<i>Reference Methods</i>									
Pre-trained	73.7	56.6	82.4	91.1	85.5	85.6	94.3	87.5	82.1
Individual	80.2	88.5	89.2	94.4	87.2	91.7	95.2	90.9	89.6
<i>Multi-Task Model Fusion Methods</i>									
Weight Averaging	74.6	84.3	84.1	92.8	86.3	87.4	94.8	88.0	86.5
Task Arithmetic	76.9	85.4	85.3	93.9	85.8	88.1	95.2	87.8	87.3
Ties-Merging	77.1	85.1	86.3	93.9	86.0	87.7	95.1	88.0	87.4
Layer-wise AdaMerging	76.7	87.6	84.8	93.8	85.9	88.1	95.2	88.6	87.6

From the experimental results in Section E.4.1 and E.4.2, we have the following key observations: (1) The performance of all multi-task model fusion methods on unseen tasks is generally lower than their performance on seen tasks. This is expected, as the models being fused are not explicitly trained on the unseen tasks. (2) A *negative transfer* is observed in Table 15 on the RESISC45 dataset, where the merged models exhibit lower accuracy compared to the pre-trained model. The performance of all multi-task model fusion methods on RESISC45 is lower than the pre-trained model, indicating that the knowledge transferred from the seen tasks may not be beneficial or even harmful to this specific unseen task. (3) The performance of all methods drops significantly on certain types of corruptions, such as pixelation and impulse noise. This highlights the challenge of maintaining robustness under severe distribution shifts and the need for further research in this direction. (4) When the test distribution is corrupted, adaptive methods may overfit to certain tasks, leading to a

decrease in overall performance. This suggests that adaptive methods may need to be further regularized to improve generalization and robustness.

E.4.1 GENERALIZATION EXPERIMENTS

Table 14: Generalization results when merging ViT-B/32 models on six tasks (SUN397, Cars, RESISC45, DTD, SVHN, GTSRB) and evaluating on two unseen tasks (MNIST, EuroSAT).

Method	Seen Tasks (ACC)							Unseen Tasks (ACC)		
	SUN397	Cars	RESISC45	DTD	SVHN	GTSRB	Avg.	MNIST	EuroSAT	Avg.
Pre-trained	63.2	59.9	60.6	43.9	23.5	30.4	46.9	47.6	45.6	46.6
Fisher Merging	65.5	67.2	78.2	57.6	84.2	75.9	71.4	71.8	49.4	60.6
RegMean	68.7	70.0	86.5	65.9	93.9	86.7	78.6	82.2	49.3	65.7
Task Arithmetic	64.3	63.0	73.2	54.9	84.7	79.5	69.9	75.5	42.6	59.1
Ties-Merging	68.3	65.5	76.9	54.9	75.4	72.0	68.9	73.1	47.3	60.2
AdaMerging	68.4	71.9	87.9	69.1	92.2	93.8	80.5	77.7	47.3	62.5
<u>WEMoE</u>	75.4	77.5	94.3	77.0	96.8	98.7	86.6	78.3	44.0	61.1

Table 15: Generalization results when merging ViT-B/32 models on six tasks (SUN397, Cars, GTSRB, EuroSAT, DTD, MNIST) and evaluating on two unseen tasks (RESISC45, SVHN).

Method	Seen Tasks (ACC)							Unseen Tasks (ACC)		
	SUN397	Cars	GTSRB	EuroSAT	DTD	MNIST	Avg.	RESISC45	SVHN	Avg.
Pre-trained	63.2	59.9	30.4	45.6	43.9	47.6	48.4	60.6	23.5	40.1
Fisher Merging	68.1	67.4	67.2	86.4	58.6	81.6	71.5	60.2	42.5	51.3
RegMean	69.4	70.5	86.9	97.0	67.1	98.3	81.5	50.2	51.5	50.8
Task Arithmetic	65.2	63.6	76.1	87.1	56.4	94.2	73.8	52.4	45.2	48.8
Ties-Merging	68.2	65.9	70.0	81.2	56.0	89.0	71.7	60.3	47.3	53.8
AdaMerging	69.8	72.4	95.5	95.1	70.7	98.1	83.6	48.7	60.7	54.7
<u>WEMoE</u>	74.3	78.1	98.8	98.7	75.1	99.5	87.4	47.3	51.3	49.3

For the generalization experiments, we assess the performance of multi-task deep model fusion algorithms on two unseen tasks after merging ViT-B/32 models trained on six tasks. The performance of various multi-task model fusion methods, including Fisher Merging (Matena and Raffel, 2022), RegMean (Jin et al., 2023), Task Arithmetic (Ilharco et al., 2023), Ties-Merging (Yadav et al., 2023), AdaMerging (Yang et al., 2024b), and WEMoE (Tang et al., 2024b), is compared across both the seen tasks and unseen tasks. Specifically, we conduct two sets of generalization experiments using the CLIP-ViT-B/32 models:

- In the first set, we merge models trained on six tasks (SUN397, Cars, RESISC45, DTD, SVHN, GTSRB) and evaluate the fused model on the unseen tasks (MNIST, EuroSAT). The results are shown in Table 14.
- In the second set of experiments, we merge models trained on six tasks (SUN397, Cars, GTSRB, EuroSAT, DTD, MNIST) and evaluate the fused model on the unseen tasks (RESISC45, SVHN). The results are shown in Table 15.

Tables 14 and 15 present the generalization performance of various multi-task model fusion algorithms when merging CLIP-ViT-B/32 models trained on six seen tasks and evaluating their performance on two unseen tasks. This analysis helps us understand how well the fused models can adapt to new tasks that were not encountered during the training and model fusion process.

E.4.2 ROBUSTNESS EXPERIMENTS

Table 16: Ablations of the test data distribution on ViT-B/32 (for all methods, $\lambda = 0.3$). Here underlined method names indicate methods that require test-time adaptation training, while *italicized method names* indicate model mixing approaches.

Method	Cars	EuroSAT	RESISC45	GTSRB	Avg.	Cars	EuroSAT	RESISC45	GTSRB	Avg.
	Clean Test Set					Corrupted Test Set (Motion Blur)				
Fisher Merging	66.0	92.7	83.7	78.7	80.3	60.7	57.6	81.7	78.4	69.6
RegMean	72.1	97.5	88.9	93.9	88.1	70.0	71.3	87.5	86.8	78.9
Task Arithmetic	64.6	91.8	80.2	74.8	77.9	62.4	59.2	78.5	63.3	65.9
Ties-Merging	65.2	83.3	78.1	67.4	73.5	64.4	53.9	76.4	57.1	62.9
AdaMerging	75.2	94.3	87.6	96.7	88.5	72.4	72.7	85.3	94.3	81.2
<u>WEMoE</u>	77.4	98.9	94.4	99.0	92.4	76.5	74.2	93.7	97.4	85.5
	Corrupted Test Set (Impulse Noise)					Corrupted Test Set (Gaussian Noise)				
Fisher Merging	61.5	50.0	74.7	52.6	59.7	61.6	48.1	76.0	51.3	59.3
RegMean	66.9	51.0	80.6	68.7	66.8	69.4	41.8	84.0	67.7	65.7
Task Arithmetic	59.8	53.3	72.3	45.0	57.6	61.5	52.5	75.0	50.1	59.8
Ties-Merging	60.2	45.6	69.8	38.3	53.5	61.8	47.3	73.1	42.3	56.1
AdaMerging	69.2	40.0	79.6	83.3	68.0	70.0	53.3	82.1	80.0	71.4
<u>WEMoE</u>	75.1	9.7	91.5	91.8	67.0	76.5	9.6	92.7	88.7	66.8
	Corrupted Test Set (Pixelate)					Corrupted Test Set (Spatter)				
Fisher Merging	2.2	34.0	17.0	63.2	29.1	61.4	64.2	74.6	47.3	61.9
RegMean	2.3	38.3	18.2	89.4	37.0	67.7	60.0	81.3	81.9	72.7
Task Arithmetic	2.3	33.2	19.1	65.6	30.0	61.0	62.5	72.8	57.0	63.3
Ties-Merging	3.3	31.8	18.0	58.5	27.9	61.3	52.9	70.3	48.1	58.2
AdaMerging	1.3	52.9	21.0	91.0	41.5	68.4	55.9	78.3	92.3	73.7
<u>WEMoE</u>	0.5	11.6	2.3	97.5	28.0	75.1	9.7	91.4	96.3	68.1
	Corrupted Test Set (Contrast)					Corrupted Test Set (JPEG Compression)				
Fisher Merging	63.8	58.4	75.5	70.4	67.0	66.3	67.6	82.6	58.9	68.8
RegMean	69.6	64.8	84.4	90.0	77.2	71.5	72.6	88.7	82.2	78.7
Task Arithmetic	62.3	55.7	75.3	70.8	66.0	63.9	66.1	80.1	61.0	67.8
Ties-Merging	64.2	52.4	74.8	63.5	63.7	65.0	59.5	77.9	53.2	63.9
AdaMerging	73.1	67.4	83.0	96.2	79.9	72.9	70.7	86.3	90.6	80.1
<u>WEMoE</u>	77.2	34.7	93.1	98.4	75.9	77.3	61.0	94.1	95.7	82.0

To assess the robustness and real-world applicability of multi-task model fusion algorithms, we conduct evaluations on corrupted test sets as presented in Table 16. These experiments

are crucial for understanding how fusion methods perform under distribution shifts and data quality degradation, which are common challenges in practical deployment scenarios. We evaluate six different corruption types across four representative image classification tasks (Cars, EuroSAT, RESISC45, GTSRB), including motion blur, impulse noise, Gaussian noise, pixelation, spatter, and JPEG compression.

The results reveal varying degrees of robustness across different fusion methods, with some approaches maintaining reasonable performance under mild corruptions while experiencing significant degradation under severe distortions such as pixelation. Notably, test-time adaptation-based methods like AdaMerging and WEMoE demonstrate better resilience compared to simpler averaging techniques, though all methods show substantial performance drops under extreme corruption conditions, highlighting the ongoing challenge of maintaining model robustness in multi-task fusion scenarios.

E.5 Scaling to Large-Scale Neural Networks

Table 17: Comparison of individual Mistral-7B models and the upscaled model on various benchmark tasks. For our method, we set $k_{gate} = 8$, $k = 512$, and the total parameter count is 11.2B. For a better comparison, we also include the performance of the Qwen1.5-14B model as a reference.

Model	MMLU	TruthfulQA	GSM8K	ARC	Challenge
Mistral-7B-v0.1 (pre-trained)	59.64	42.62	38.81		53.92
Qwen1.5-14B (reference)	66.11	52.00	69.37		49.93
MetaMath-Mistral-7B	60.56	44.79	71.49		51.02
dolphin-2.1-mistral-7b	60.56	55.88	56.93		57.00
speechless-code-mistral-7b-v1.0	61.18	47.47	48.98		<u>57.68</u>
Simple Average	61.42	49.95	67.40		57.59
Task Arithmetic ($\lambda = 0.3$)	<u>61.29</u>	49.38	66.94		57.94
SMILE Upscaled model ($k = 512$, 11.2B)	60.66	52.79	67.85		54.35
SMILE Upscaled model (Dense experts)	60.61	<u>54.23</u>	<u>70.66</u>		55.12

Model fusion methods can also be applied to large-scale neural networks including Large Language Models (LLMs) and Multimodal Large Language Models (MLLMs). The high computational cost associated with developing LLMs are a significant practical challenge for many researchers. In FusionBench, we have developed multiple model fusion techniques applicable to LLMs for cheap and efficient model scaling, which can save computational resources for developing new models. Moreover, FusionBench is actively developing and integrating more LLM and MLLM merging methods, with ongoing work to expand benchmark coverage to large-scale model fusion tasks. While evaluation methodologies for LLMs and MLLMs are still maturing and computational costs remain a challenge, FusionBench supports established evaluation frameworks such as LM-Evaluation-Harness (Gao et al., 2024) for assessing model performance. As the toolkit evolves, more comprehensive LLM task evaluations will be included in future releases.

Take merging Mistral-7B models using SMILE upscaling (Tang et al., 2025a) as an example, we compare the performance of the individual Mistral-7B models and the upscaled

model on various benchmark tasks in Table 17, as well as the performance of the Qwen1.5-14B model as a reference. The LM-Evaluation-Harness is utilized to assess the performance of the models. We fuse three Mistral-7B models, each fine-tuned for a distinct downstream task as showing varying performance across different tasks, thereby incorporating task-specific expertise. It is observed that the upscaled models and merged methods (Simple Average, Task Arithmetic) generally enhance performance compared to individual models, demonstrating the benefits of model fusion techniques.

Appendix F. Prompt-Based Text-to-Text Generation

This section details the prompt templates employed for each of the eight text-to-text generation tasks from the GLUE benchmark, see Section E.3 for more details. Within each task, we provide the format of the input text, and the corresponding target text mapping. These templates are crucial in fine-tuning the Flan-T5 models for generating appropriate text outputs tailored to each specific task.

- CoLA:
 - *Input Text*: "Indicate if the following sentence is grammatically correct or not: "sentence". Answer 'acceptable' or 'unacceptable'."
 - *Target Text*:
 - * 0: "unacceptable"
 - * 1: "acceptable"
- MNLI:
 - *Input Text*: "Does the premise: 'premise' logically imply, contradict, or is neutral to the hypothesis: 'hypothesis'? Answer with 'entailment', 'contradiction', or 'neutral'."
 - *Target Text*:
 - * 0: "entailment"
 - * 1: "neutral"
 - * 2: "contradiction"
- MRPC:
 - *Input Text*: "Are the following sentences 'sentence1' and 'sentence2' conveying the same meaning? Answer with 'yes' or 'no'."
 - *Target Text*:
 - * 0: "no"
 - * 1: "yes"
- QNLI:
 - *Input Text*: "Given the context: 'sentence', does the question 'question' have an answer based on the information provided? Answer with 'yes' or 'no'."

- *Target Text*:
 - * 0: "yes"
 - * 1: "no"
- QQP:
 - *Input Text*: "Do the questions ‘question1’ and ‘question2’ have the same intent? Answer with ‘yes’ or ‘no’."
 - *Target Text*:
 - * 0: "no"
 - * 1: "yes"
- RTE:
 - *Input Text*: "Does the text: ‘sentence1’ entail that ‘sentence2’ is true? Provide ‘yes’ or ‘no’."
 - *Target Text*:
 - * 0: "yes"
 - * 1: "no"
- SST-2:
 - *Input Text*: "Given the sentence ‘sentence’, determine the sentiment. Is it positive or negative?"
 - *Target Text*:
 - * 0: "negative"
 - * 1: "positive"
- STSB:
 - *Input Text*: "Consider the sentences ‘sentence1’ and ‘sentence2’. On a scale from 1 (completely different) to 5 (completely similar), rate the similarity."
 - *Target Text*: ".:1f", parse to float with one decimal place

Reporting Metrics: We report accuracy for all tasks except for STSB, where we use Spearman’s ρ as the evaluation metric. For task STSB, the model is expected to output a numerical value. An example from the STSB task is as follows:

- *Input*:
 - Sentence 1: A plane is taking off.
 - Sentence 2: An air plane is taking off.
- *Output*:
 - label: 5

We try to parse the output as a numerical value. If the model outputs a numerical value, we can calculate the Spearman’s rho between the predicted numerical value and the ground truth numerical value. If the model outputs a non-numerical value, we assume the Spearman’s rho is 0, indicating that there is no discernible monotonic increasing or decreasing relationship between the model’s predictions and the ground truth. This is a conservative approach, as even non-numerical outputs might contain some relevant information that’s being discarded in this evaluation.

References

- Samuel Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git Re-Basin: Merging Models modulo Permutation Symmetries. 2023. URL <https://openreview.net/forum?id=CQsmMYmlP5T>.
- James Urquhart Allingham, Florian Wenzel, Zelda E Mariet, Basil Mustafa, Joan Puigcerver, Neil Houlsby, Ghassen Jerfel, Vincent Fortuin, Balaji Lakshminarayanan, Jasper Snoek, et al. Sparse moes meet efficient ensembles. *arXiv preprint arXiv:2110.03360*, 2021.
- Umar Asif, Jianbin Tang, and Stefan Herrer. Ensemble knowledge distillation for learning improved and efficient networks. In *ECAI 2020*, pages 953–960. IOS Press, 2020.
- Gregory Benton, Wesley Maddox, Sanae Lotfi, and Andrew Gordon Gordon Wilson. Loss surface simplexes for mode connecting volumes and fast ensembling. In *International Conference on Machine Learning*, pages 769–779. PMLR, 2021.
- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101–mining discriminative components with random forests. In *European conference on computer vision*, pages 446–461. Springer, 2014.
- Anthony Chen, Huanrui Yang, Yulu Gan, Denis A. Gudovskiy, Zhen Dong, Haofan Wang, Tomoyuki Okuno, Yohei Nakata, Kurt Keutzer, and Shanghang Zhang. Split-Ensemble: Efficient OOD-aware Ensemble via Task and Model Splitting. pages 7568–7585, July 2024. ISSN 2640-3498. URL <https://proceedings.mlr.press/v235/chen24aw.html>.
- Hao Mark Chen, Shell Xu Hu, Wayne Luk, Timothy Hospedales, and Hongxiang Fan. FW-Merging: Scaling Model Merging with Frank-Wolfe Optimization. pages 3390–3400, 2025. URL https://openaccess.thecvf.com/content/ICCV2025/html/Chen_FW-Merging_Scaling_Model_Merging_with_Frank-Wolfe_Optimization_ICCV_2025_paper.html.
- Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017.
- Alexandra Chronopoulou, Matthew Peters, Alexander Fraser, and Jesse Dodge. AdapterSoup: Weight Averaging to Improve Generalization of Pretrained Language Models. pages 2054–2063, May 2023. doi: 10.18653/v1/2023.findings-eacl.153. URL <https://aclanthology.org/2023.findings-eacl.153/>.

- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014.
- Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*, 2018.
- Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011.
- Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE, 2017.
- Nico Daheim, Thomas Möllenhoff, Edoardo Ponti, Iryna Gurevych, and Mohammad Emtiyaz Khan. Model Merging by Uncertainty-Based Gradient Matching. 2024. URL <https://openreview.net/forum?id=D7KJmfEDQP>.
- Damai Dai, Chengqi Deng, Chenggang Zhao, R.x. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y.k. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models. pages 1280–1297, August 2024. doi: 10.18653/v1/2024.acl-long.70. URL <https://aclanthology.org/2024.acl-long.70/>.
- Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*, pages 1309–1318. PMLR, 2018.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pages 3259–3269. PMLR, 2020.
- Elias Frantar and Dan Alistarh. Sparsegpt: Massive language models can be accurately pruned in one-shot. In *International conference on machine learning*, pages 10323–10337. PMLR, 2023.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, 07 2024. URL <https://zenodo.org/records/12608602>.

- Antonio Andrea Gargiulo, Donato Crisostomi, Maria Sofia Bucarelli, Simone Scardapane, Fabrizio Silvestri, and Emanuele Rodolà. Task Singular Vectors: Reducing Task Interference in Model Merging. pages 18695–18705, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Gargiulo_Task_Singular_Vectors_Reducing_Task_Interference_in_Model_Merging_CVPR_2025_paper.html.
- Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018.
- Charles Goddard, Shamane Siriwardhana, Malikeh Ehghaghi, Luke Meyers, Vlad Karpukhin, Brian Benedict, Mark McQuade, and Jacob Solawetz. Arcee’s mergekit: A toolkit for merging large language models. *arXiv preprint arXiv:2403.13257*, 2024.
- Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *International conference on neural information processing*, pages 117–124. Springer, 2013.
- Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. 2016. URL <http://arxiv.org/abs/1510.00149>.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
- Yifei He, Siqi Zeng, Yuzheng Hu, Rui Yang, Tong Zhang, and Han Zhao. MergeBench: A Benchmark for Merging Domain-Specialized LLMs. October 2025. URL [https://openreview.net/forum?id=rw50iUoyLu&referrer=%5Bthe%20profile%20of%20Yuzheng%20Hu%5D\(%2Fprofile%3Fid%3D~Yuzheng_Hu1\)](https://openreview.net/forum?id=rw50iUoyLu&referrer=%5Bthe%20profile%20of%20Yuzheng%20Hu%5D(%2Fprofile%3Fid%3D~Yuzheng_Hu1)).
- Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 204–207. IEEE, 2018.
- Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Ming Hu, Zhihao Yue, Xiaofei Xie, Cheng Chen, Yihao Huang, Xian Wei, Xiang Lian, Yang Liu, and Mingsong Chen. Is Aggregation the Only Choice? Federated Learning via Layer-wise Model Recombination. pages 1096–1107, August 2024. doi: 10.1145/3637528.3671722. URL <https://doi.org/10.1145/3637528.3671722>.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*, 2023.

- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. Personalized Soups: Personalized Large Language Model Alignment via Post-hoc Parameter Merging. November 2024. URL <https://openreview.net/forum?id=EMrnoPRvxe>.
- Yushu Jiang. Evomerge: Neuroevolution for large language models. *arXiv preprint arXiv:2402.00070*, 2024.
- Xisen Jin, Xiang Ren, Daniel Preotiuc-Pietro, and Pengxiang Cheng. Dataless Knowledge Fusion by Merging Weights of Language Models. 2023. URL <https://openreview.net/forum?id=FCnohuR6AnM>.
- Sanghoon Kim, Dahyun Kim, Chanjun Park, Wonsung Lee, Wonho Song, Yunsu Kim, Hyeonwoo Kim, Yungi Kim, Hyeonju Lee, Jihoo Kim, Changbae Ahn, Seonghoon Yang, Sukyung Lee, Hyunbyung Park, Gyoungjin Gim, Mikyoung Cha, Hwalsuk Lee, and Sunghun Kim. SOLAR 10.7B: Scaling Large Language Models with Simple yet Effective Depth Up-Scaling. pages 23–35, June 2024. doi: 10.18653/v1/2024.naacl-industry.3. URL <https://aclanthology.org/2024.naacl-industry.3/>.
- Aran Komatsuzaki, Joan Puigcerver, James Lee-Thorp, Carlos Riquelme Ruiz, Basil Mustafa, Joshua Ainslie, Yi Tay, Mostafa Dehghani, and Neil Houlsby. Sparse Upcycling: Training Mixture-of-Experts from Dense Checkpoints. 2023. URL <https://openreview.net/forum?id=T5nUQDrM4u>.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yann Lecun, Le’on Bottou, Yoshua Bengio, and Parick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. ISSN 00189219. doi: 10.1109/5.726791. URL <http://ieeexplore.ieee.org/document/726791/>.
- Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 991–999, 2015.
- Weishi Li, Yong Peng, Miao Zhang, Liang Ding, Han Hu, and Li Shen. Deep model fusion: A survey. *arXiv preprint arXiv:2309.15698*, 2023.
- Chang Liu, Chenfei Lou, Runzhong Wang, Alan Yuhan Xi, Li Shen, and Junchi Yan. Deep neural network fusion via graph matching with applications to model ensemble and federated learning. In *International Conference on Machine Learning*, pages 13857–13869. PMLR, 2022a.

- James Liu, Guangxuan Xiao, Kai Li, Jason D Lee, Song Han, Tri Dao, and Tianle Cai. Bitdelta: Your fine-tune may only be worth one bit. *Advances in Neural Information Processing Systems*, 37:13579–13600, 2024.
- Tian Yu Liu and Stefano Soatto. Tangent model composition for ensembling and continual fine-tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18676–18686, 2023.
- Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022b.
- Xudong Lu, Qi Liu, Yuhui Xu, Aojun Zhou, Siyuan Huang, Bo Zhang, Junchi Yan, and Hongsheng Li. Not All Experts are Equal: Efficient Expert Pruning and Skipping for Mixture-of-Experts Large Language Models. pages 6159–6172, August 2024a. doi: 10.18653/v1/2024.acl-long.334. URL <https://aclanthology.org/2024.acl-long.334/>.
- Zhenyi Lu, Chenghao Fan, Wei Wei, Xiaoye Qu, Dangyang Chen, and Yu Cheng. Twin-Merging: Dynamic Integration of Modular Expertise in Model Merging. *Advances in Neural Information Processing Systems*, 37:78905–78935, December 2024b. doi: 10.52202/079017-2505. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/8fcd17eb91bae20d9826786d7d6be799-Abstract-Conference.html.
- Daniel Marczak, Simone Magistri, Sebastian Cygert, Bartłomiej Twardowski, Andrew D. Bagdanov, and Joost van de Weijer. No Task Left Behind: Isotropic Model Merging with Common and Task-Specific Subspaces. June 2025. URL <https://openreview.net/forum?id=RBZpAa271s¬eId=mUEC1Bk5g0>.
- Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. Fast Model Editing at Scale. October 2021. URL <https://openreview.net/forum?id=ODcZxeWf0Pt>.
- Luca Moschella, Valentino Maiorca, Marco Fumero, Antonio Norelli, Francesco Locatello, and Emanuele Rodolà. Relative representations enable zero-shot latent space communication. 2023. URL <https://openreview.net/forum?id=SrC-nwieGJ>.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, page 4. Granada, 2011.
- The-Hai Nguyen, Dang Huu-Tien, Takeshi Suzuki, and Le-Minh Nguyen. Regmean++: Enhancing effectiveness and generalization of regression mean for model merging. *arXiv preprint arXiv:2508.03121*, 2025.
- Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008.

- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=a68SUt6zFt>. Featured Certification.
- Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Oleksiy Ostapenko, Zhan Su, Edoardo Maria Ponti, Laurent Charlin, Nicolas Le Roux, Lucas Caccia, and Alessandro Sordoni. Towards modular LLMs by building and reusing a library of LoRAs. 235:38885–38904, July 2024.
- Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012.
- Manas Pathak, Shantanu Rane, and Bhiksha Raj. Multiparty differential privacy via aggregation of locally trained classifiers. *Advances in neural information processing systems*, 23, 2010.
- Xingyu Qu and Samuel Horvath. Rethink model re-basin and the linear mode connectivity. *arXiv preprint arXiv:2402.05966*, 2024.
- Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor, Laure Soulier, and Matthieu Cord. Rewarded soups: towards pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards. *Advances in Neural Information Processing Systems*, 36, 2024.
- Omer Sagi and Lior Rokach. Ensemble learning: A survey. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 8(4):e1249, 2018.
- Li Shen, Anke Tang, Enneng Yang, Guibing Guo, Yong Luo, Lefei Zhang, Xiaochun Cao, Bo Du, and Dacheng Tao. Efficient and Effective Weight-Ensembling Mixture of Experts for Multi-Task Model Merging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (01):1–13, November 2025. ISSN 0162-8828. doi: 10.1109/TPAMI.2025.3629605. URL <https://www.computer.org/csdl/journal/tp/5555/01/11230232/2bqyGYH5Hby>.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- Johannes Stalkamp, Marc Schlippsing, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012.

- George Stoica, Daniel Bolya, Jakob Brandt Bjorner, Pratik Ramesh, Taylor Hearn, and Judy Hoffman. ZipIt! Merging Models from Different Tasks without Training. October 2023. URL <https://openreview.net/forum?id=LEYUkvdUhq>.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. A Simple and Effective Pruning Approach for Large Language Models. 2024. URL <https://openreview.net/forum?id=PxoFut3dWW>.
- Derek Tam, Mohit Bansal, and Colin Raffel. Merging by matching models in task parameter subspaces. *Transactions on Machine Learning Research*, 2024.
- Anke Tang, Yong Luo, Han Hu, Fengxiang He, Kehua Su, Bo Du, Yixin Chen, and Dacheng Tao. Improving Heterogeneous Model Reuse by Density Estimation. In *Thirty-Second International Joint Conference on Artificial Intelligence*, volume 4, pages 4244–4252, August 2023a. doi: 10.24963/ijcai.2023/472. URL <https://www.ijcai.org/proceedings/2023/472>.
- Anke Tang, Li Shen, Yong Luo, Liang Ding, Han Hu, Bo Du, and Dacheng Tao. Concrete subspace learning based interference elimination for multi-task model fusion. *arXiv preprint arXiv:2312.06173*, 2023b.
- Anke Tang, Li Shen, Yong Luo, Shiwei Liu, Han Hu, and Bo Du. Towards efficient pareto set approximation via mixture of experts based model fusion. *arXiv preprint arXiv:2406.09770*, 2024a.
- Anke Tang, Li Shen, Yong Luo, Nan Yin, Lefei Zhang, and Dacheng Tao. Merging Multi-Task Models via Weight-Ensembling Mixture of Experts. pages 47778–47799, July 2024b. ISSN 2640-3498. URL <https://proceedings.mlr.press/v235/tang24e.html>.
- Anke Tang, Li Shen, Yong Luo, Shuai Xie, Han Hu, Lefei Zhang, Bo Du, and Dacheng Tao. Zero-Shot Sparse Mixture of Low-Rank Experts Construction From Pre-Trained Foundation Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–12, 2025a. ISSN 1939-3539. doi: 10.1109/TPAMI.2025.3612480. URL <https://ieeexplore.ieee.org/document/11175033>.
- Anke Tang, Enneng Yang, Li Shen, Yong Luo, Han Hu, Lefei Zhang, Bo Du, and Dacheng Tao. Merging on the Fly Without Retraining: A Sequential Approach to Scalable Continual Model Merging. October 2025b. URL [https://openreview.net/forum?id=rdGMyTPhui&referrer=%5Bthe%20profile%20of%20Bo%20Du%5D\(%2Fprofile%3Fid%3D~Bo_Du3\)](https://openreview.net/forum?id=rdGMyTPhui&referrer=%5Bthe%20profile%20of%20Bo%20Du%5D(%2Fprofile%3Fid%3D~Bo_Du3)).
- Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. In *International Conference on Medical image computing and computer-assisted intervention*, pages 210–218. Springer, 2018.
- Fanqi Wan, Xinting Huang, Deng Cai, Xiaojun Quan, Wei Bi, and Shuming Shi. Knowledge Fusion of Large Language Models. October 2023. URL <https://openreview.net/forum?id=jiDsk12qcz>.

- Fanqi Wan, Longguang Zhong, Ziyi Yang, Ruijun Chen, and Xiaojun Quan. FuseChat: Knowledge Fusion of Chat Models. pages 21629–21653, November 2025. doi: 10.18653/v1/2025.emnlp-main.1096. URL <https://aclanthology.org/2025.emnlp-main.1096/>.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. pages 353–355, November 2018. doi: 10.18653/v1/W18-5446. URL <https://aclanthology.org/W18-5446/>.
- Ke Wang, Nikolaos Dimitriadis, Guillermo Ortiz-Jiménez, François Fleuret, and Pascal Frossard. Localizing task information for improved model merging and compression. In *Proceedings of the 41st International Conference on Machine Learning*, pages 50268–50287, 2024.
- Yeming Wen, Dustin Tran, and Jimmy Ba. BatchEnsemble: An Alternative Approach to Efficient Ensemble and Lifelong Learning. September 2019. URL <https://openreview.net/forum?id=Sk1f1yrYDr>.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR, 2022.
- Xi-Zhu Wu, Song Liu, and Zhi-Hua Zhou. Heterogeneous model reuse via optimizing multiparty multiclass margin. In *International Conference on Machine Learning*, pages 6840–6849. PMLR, 2019.
- Xi-Zhu Wu, Wenkai Xu, Song Liu, and Zhi-Hua Zhou. Model reuse with reduced kernel mean embedding specification. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):699–710, 2021.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010.
- Yichu Xu, Xin-Chun Li, Le Gan, and De-Chuan Zhan. Weight scope alignment: A frustratingly easy method for model merging. *arXiv preprint arXiv:2408.12237*, 2024.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin A. Raffel, and Mohit Bansal. TIES-Merging: Resolving Interference When Merging Models. *Advances in Neural Information Processing Systems*, 36:7093–7115, December 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1644c9af28ab7916874f6fd6228a9bcf-Abstract-Conference.html.

- Prateek Yadav, Colin Raffel, Mohammed Muqeeth, Lucas Caccia, Haokun Liu, Tianlong Chen, Mohit Bansal, Leshem Choshen, and Alessandro Sordoni. A Survey on Model MoErging: Recycling and Routing Among Specialized Experts for Collaborative Learning. *Transactions on Machine Learning Research*, August 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=u0azVc9Y0y>.
- Prateek Yadav, Tu Vu, Jonathan Lai, Alexandra Chronopoulou, Manaal Faruqui, Mohit Bansal, and Tsendsuren Munkhdalai. What Matters for Model Merging at Scale? *Transactions on Machine Learning Research*, May 2025. ISSN 2835-8856. URL <https://openreview.net/forum?id=9sbtmVnpW>.
- Enneng Yang, Li Shen, Zhenyi Wang, Guibing Guo, Xiaojun Chen, Xingwei Wang, and Dacheng Tao. Representation surgery for multi-task model merging. *Forty-first International Conference on Machine Learning*, 2024a.
- Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. *The Twelfth International Conference on Learning Representations*, 2024b.
- Hanrong Ye and Dan Xu. Taskexpert: Dynamically assembling multi-task representations with memorial mixture-of-experts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21828–21837, 2023.
- Xin Yi, Shunfan Zheng, Linlin Wang, Xiaoling Wang, and Liang He. A safety realignment framework via subspace-oriented model fusion for large language models. *Knowledge-Based Systems*, 306:112701, December 2024. ISSN 0950-7051. doi: 10.1016/j.knosys.2024.112701. URL <https://www.sciencedirect.com/science/article/pii/S0950705124013352>.
- Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*, 2024.
- Ningyu Zhang, Yunzhi Yao, Bozhong Tian, Peng Wang, Shumin Deng, Mengru Wang, Zekun Xi, Shengyu Mao, Jintian Zhang, Yuansheng Ni, et al. A comprehensive study of knowledge editing for large language models. *arXiv preprint arXiv:2401.01286*, 2024.
- Ziyu Zhao, Leilei Gan, Guoyin Wang, Wangchunshu Zhou, Hongxia Yang, Kun Kuang, and Fei Wu. LoraRetriever: Input-Aware LoRA Retrieval and Composition for Mixed Tasks in the Wild. pages 4447–4462, August 2024. doi: 10.18653/v1/2024.findings-acl.263. URL <https://aclanthology.org/2024.findings-acl.263/>.
- Hongling Zheng, Li Shen, Anke Tang, Yong Luo, Han Hu, Bo Du, Yonggang Wen, and Dacheng Tao. Learning from models beyond fine-tuning. *Nature Machine Intelligence*, pages 1–12, 2025.
- Hangyu Zhou, Aaron Gokaslan, Volodymyr Kuleshov, and Bharath Hariharan. Randes: Randomized delta superposition for multi-model compression. *arXiv preprint arXiv:2505.11204*, 2025.

Zhi Hua Zhou. Learnware: On the future of machine learning. *Frontiers of Computer Science*, 10(4):589–590, 2016. ISSN 20952236. doi: 10.1007/s11704-016-6906-3.

Didi Zhu, Zhongyi Sun, Zexi Li, Tao Shen, Ke Yan, Shouhong Ding, Chao Wu, and Kun Kuang. Model tailor: Mitigating catastrophic forgetting in multi-modal large language models. 235:62581–62598, July 2024.