

Med-LEGO: Editing and Adapting toward Generalist Medical Image Diagnosis

Yitao Zhu¹, Yuan Yin¹, Jiaming Li¹, Mengjie Xu¹, Zihao Zhao¹,
Honglin Xiong¹, Sheng Wang¹, and Qian Wang^{1, 2}(✉)

¹ School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai, 201210, China

² Shanghai Clinical Research and Trial Center, Shanghai, 201210, China
qianwang@shanghaitech.edu.cn

Abstract. The adoption of visual foundation models has become a common practice in computer-aided diagnosis (CAD). While these foundation models provide a viable solution for creating generalist medical AI, privacy concerns make it difficult to pre-train or continuously update such models across multiple domains and datasets, leading many studies to focus on specialist models. To address this challenge, we propose Med-LEGO, a training-free framework that enables the seamless integration or updating of a generalist CAD model by combining multiple specialist models, similar to assembling LEGO bricks. Med-LEGO enhances LoRA (low-rank adaptation) by incorporating singular value decomposition (SVD) to efficiently capture the domain expertise of each specialist model with minimal additional parameters. By combining these adapted weights through simple operations, Med-LEGO allows for the easy integration or modification of specific diagnostic capabilities without the need for original data or retraining. Finally, the combined model can be further adapted to new diagnostic tasks, making it a versatile generalist model. Our extensive experiments demonstrate that Med-LEGO outperforms existing methods in both cross-domain and in-domain medical tasks while using only 0.18% of full model parameters. These merged models show better convergence and generalization to new tasks, providing an effective path toward generalist medical AI.

Keywords: Model Editing · Medical Image Diagnosis · Generalist AI · Knowledge Decomposition · Parameter Efficient Fine-Tuning.

1 Introduction

Visual foundation models have gained significant attention from the medical image diagnosis community [10], serving as the backbone for many computer-aided diagnosis (CAD) systems. Vision Transformers [2] and similar models are commonly used in this context. However, developing medical foundation models is challenging due to limited annotated data, modality variations, and high data acquisition costs [11]. As a result, fine-tuning pre-trained models on small,

task-specific datasets has become a common and practical approach in data-limited medical imaging [13]. However, this approach faces three key challenges: (1) Task-specific models hinder generalization: Each new task requires a separate model, making it difficult to develop a unified, general-purpose model. This fragmentation limits scalability and adaptability. (2) Weak inter-task relationships: Models trained in isolation cannot leverage knowledge from related tasks, restricting both in-domain performance and out-of-domain generalization. This limits the potential for cross-task learning. (3) The computational costs of continuously updating foundation models for new tasks and data pose significant practical challenges. These issues highlight the need for more efficient and integrated approaches in medical imaging.

To address the limitations of isolated task-specific models, recent progress in model editing has introduced efficient ways to assemble pre-trained models without extensive retraining. These methods aim to integrate knowledge from multiple tasks into a unified model, enabling improvements in generalization, downstream task performance, bias mitigation, and the incorporation of new information. A key advancement in this area is Task Arithmetic [4], which introduces task vectors—compact representations derived by subtracting pre-trained weights from task-specific fine-tuned weights. These vectors allow models to adapt to new tasks through simple arithmetic operations, such as adding vectors to enhance multi-task performance or negating them to remove undesired behaviors. However, Task Arithmetic relies on manually defined task correlations, and directly performing addition or subtraction on model weights can introduce significant noise and error. For example, when two tasks that are opposites in the task vector space are added, the resulting model may perform poorly on both tasks. To address these limitations, methods like Tie-Merging [16], DARE [18], PEM Composition [19] have introduced more refined strategies, such as parameter-wise sign selection and dynamic weight adjustments, improving multi-task performance. Despite these advances, existing methods still face two key challenges: (1) Most of them operate on entire model parameters, resulting in high computational costs, and (2) they require careful design of parameter-wise strategies for sign and scale allocation during integration.

To overcome these limitations, we propose a novel approach called MedLEGO based on Singular Value Decomposition (SVD) and Low-Rank Adaptation (LoRA) [3], which efficiently captures task-specific capabilities while significantly reducing storage and computational costs compared to traditional methods. By leveraging the inherent structure of SVD, we decompose task-specific adaptations and retain only the most informative components, effectively mitigating noise and enhancing the robustness of the merged model. This approach eliminates the need for manual tuning of task correlations or scaling, as it automatically focuses on the essential shared features across tasks. The result is a more stable and generalizable model, particularly beneficial in complex domains like medical image diagnosis, where tasks are often heterogeneous and require nuanced handling. Our method not only simplifies the fusion process




Fig. 1. Overview of the Med-LEGO method: (a) We add a trainable SVD-LoRA structure to the Q and V matrices of the ViT during fine-tuning, while freezing other parameters, to speed up training and reduce the number of parameters. (b) We obtain SVD-LoRAs for different medical image datasets while sharing the same pre-trained model. (c) We combine and perform SVD decomposition on SVD-LoRAs cross-domain to extract general information for a general doctor. (d) We combine and perform SVD decomposition on SVD-LoRAs of the same domain to create a specialized doctor.

but also ensures superior multi-task performance by preserving the most critical information across diverse tasks.

The main contributions of this paper are:

- We propose Med-LEGO, a **training-free** framework that enables seamless integration of specialist models into a generalist medical AI system **without requiring access to original training data**.
- We develop SVD-LoRA, which combines SVD decomposition with Low-Rank Adaptation to efficiently capture domain expertise with **only 0.18%** of full model parameters while effectively mitigating noise during model fusion.
- Med-LEGO demonstrates **state-of-the-art performance** across 7 merging cross-domain and 3 merging in-domain medical image datasets, achieving superior generalization on 3 new medical tasks.

2 Method

2.1 Problem Statement

Given a set of N medical image diagnosis tasks $\{T_1, T_2, \dots, T_N\}$, each with its corresponding adapted model $\{M_1, M_2, \dots, M_N\}$ derived from M_{pre} , which stands for the weight of pre-training model. Our research objective is to develop

an efficient methodology to merge these task-specific adapted models into a unified model M_{merge} . The challenge lies in preserving the specialized knowledge while enhancing cross-task generalization without additional training. Moreover, for a new, related medical image diagnosis task T_{new} , using M_{merge} as the initialization weight should lead to a more stable training process and better convergence performance compared to M_{pre} .

2.2 Task Adaptation via SVD-LoRA

Directly averaging fine-tuned model weights has been shown to effectively integrate model capabilities [5, 9, 15], but for large foundation models, saving, transmitting, and computing the entire parameter space is a challenge. Meanwhile, parameter-efficient fine-tuning (PEFT) has gained popularity recently, with Low-Rank Adaptation (LoRA) [3] becoming widely adopted due to its effectiveness and efficiency. Related work in the medical image field also demonstrates its effectiveness [21]. LoRA achieves this by incrementally updating the pre-trained weights through the product of two small matrices. For the forward pass, given $h = Wx$, the modified model can be expressed as:

$$h = Wx + \Delta x = Wx + BAx, \quad (1)$$

where $W, \Delta \in \mathbb{R}^{d_m \times d_n}$, $A \in \mathbb{R}^{r \times d_n}$, and $B \in \mathbb{R}^{d_m \times r}$, with $r \ll \{d_m, d_n\}$. The rank r is kept much smaller than the dimensions d_m and d_n , which reduces the trainable number of parameters in the model.

However, merging LoRA models presents a fundamental mathematical challenge. When attempting to merge multiple LoRA models, we encounter the following inequality:

$$\underbrace{\frac{1}{N}(B_1 + \dots + B_N) \times \frac{1}{N}(A_1 + \dots + A_N)}_{\text{pre-merge}} \neq \underbrace{\frac{1}{N}(B_1 A_1 + \dots + B_N A_N)}_{\text{post-merge}}. \quad (2)$$

If we choose the pre-merge approach, although it retains the low-rank structure for easier subsequent fine-tuning, this operation violates the original rules of BAx . On the other hand, if we choose the post-merge approach, while it adheres to the original operational principles, it loses the low-rank structure and reverts to the size of $\Delta \in \mathbb{R}^{d_m \times d_n}$.

To overcome the merging challenge, and inspired by AdaLoRA [20], we propose to use SVD-LoRA instead of LoRA which is shown in Fig 1 (a), we add the SVD-LoRA structure in the Q and V weight matrix in transformer blocks. Our approach modifies the update matrix to incorporate SVD structure, expressing it as:

$$h = Wx + \Delta x = Wx + BEAx, \quad (3)$$

where $B \in \mathbb{R}^{d_m \times r}$ and $A \in \mathbb{R}^{r \times d_n}$ represent the right and left singular vectors of Δ respectively, and $E \in \mathbb{R}^{r \times r}$ is a diagonal matrix of singular values. E is initialized to zero, while B and A are initialized with random Gaussian distributions, ensuring $\Delta = 0$ at the start of training.

To maintain the orthogonality of B and A ($B^\top B = AA^\top = I$), we introduce a regularization term:

$$\mathcal{L}_{reg}(B, A) = \|B^\top B - I\|_F^2 + \|AA^\top - I\|_F^2. \quad (4)$$

This SVD structure provides two key advantages: (1) It allows for more stable training through explicit control of singular values; (2) It enables effective merging of models with different ranks through SVD decomposition explained in the next section.

2.3 Assembly of Model Capabilities

An important advantage of using SVD-LoRA as a capability representation is that it enables the fusion of parameter matrices from different ranks when using the post-merge approach. After averaging, we can easily apply the SVD decomposition to restore the resulting matrix to its original low-rank form without additional impact. The process is shown in Fig 1 (a) and (b). Specifically, for the N sets of SVD-LoRA models $\{\Delta_1, \Delta_2, \dots, \Delta_N\}$ obtained from fine-tuning on different tasks, we can merge and decompose them as follows:

$$\Delta_{\text{full}} = \hat{B}\hat{E}\hat{A} = \text{SVD}\left(\frac{1}{N}\sum_{i=1}^N B_i E_i A_i\right), \quad (5)$$

$$\Delta_{\text{merge}} = \hat{B}_{:,1:k} \hat{E}_{1:k} \hat{A}_{1:k,:}, \quad (6)$$

where the top- k values are selected based on the cumulative sum of the singular values in $\hat{E} = [e_1, e_2, \dots, e_N]$ such that the sum of the first k singular values exceeds a threshold v . This ensures that the most significant components are retained in the merged model. Finally, we obtain an updated general SVD-LoRA weight, which can be used for adaptation on new medical image datasets.

3 Experiments

3.1 Datasets

All datasets are divided into training, validation, and test sets.

- **MedMNIST:** We use 7 datasets from MedMNIST [17], including Blood (Blood Cell Microscope modality), Breast (Breast Ultrasound modality), Derma (Dermatoscope modality), Organ (Abdominal CT modality), Pathology (Colon Pathology modality), Pneumonia (Chest X-Ray modality), and Retina (Fundus Camera modality). All images are of size 224×224 .
- **NIH-CXR14:** The NIH-CXR14 dataset [14] comprises 112,120 X-ray images with disease labels from 30,805 unique patients. Each chest X-ray contains 14 binary labels for thoracic diseases.

- **Tuberculosis:** Tuberculosis was collected by Shenzhen No.3 Hospital in Shenzhen, China. It consists of 326 normal CXR images and 336 abnormal CXR images showing various manifestations of tuberculosis.
- **OAI:** OAI is a multi-center, longitudinal study on osteoarthritis, containing X-ray images of the five stages of knee osteoarthritis, classified by the Kellgren and Lawrence grading system [6].
- **Blood-Cell:** contains 12,500 augmented images of four blood cell subtypes including Eosinophil, Lymphocyte, Monocyte, and Neutrophil. Our task is to identify their blood cell types.

3.2 Implementation Details

In this paper, all images are resized to 224×224 . We use the ViT-base-patch16 model pre-trained on ImageNet [1, 12] as pre-trained weight for all ViT-based and LoRA-based methods, while PMC-CLIP [7] uses its own pre-trained weights on medical data. Following the experience from the MeLo [21] work, we set $r = 4$ for all low-rank structures. For all model training, we use cross-entropy as the loss function and train for 100 epochs with a learning rate of $3e - 4$. At the end of each epoch, we validate the model using the validation set and only perform testing and save the checkpoint when the validation accuracy reaches its optimal value. All methods converge to their optimal performance within 100 epochs. The threshold v for Med-LEGO is 99.7%.

3.3 Performance of Cross-domain Tasks Merging

We use seven medical image datasets from different modalities and tasks to validate the effectiveness of Med-LEGO as a model editing method, with results shown in Table 1. We find that among the task adaptation methods, there is no significant difference in accuracy between ViT, LoRA, and SVD-LoRA. However, for the method using multi-task pre-training, PMC-CLIP, the performance remains poor even with multi-task fine-tuning across the seven datasets. This highlights the importance of model merging in the field of medical image diagnosis. Among the model merging methods compared, we observe that other algorithm are heavily influenced by biases in individual datasets, leading to a significant performance drop across all other tasks. This issue arises because traditional methods of directly adding or subtracting model parameters struggle to balance the capabilities across different tasks, especially in the context of the large domain gap in medical image diagnosis, resulting in a strong bias toward individual tasks.

For Med-LEGO, our test results consistently outperform other methods across most tasks, demonstrating that in the field of medical image diagnosis, our SVD-LoRA representation and the fusion strategy based on SVD decomposition effectively overcome the significant differences between tasks, enabling successful cross-domain information integration.

Table 1. Comparison of the accuracy of model merging methods across 7 datasets from MedMNIST, where the gray section represents fine-tuning for each task individually. The yellow section represents multi-task fine-tuning of a single model across all tasks. The remaining section represents testing a single merged model on all tasks. **Bold** indicates the best result, and the underlined number denotes the second-best result.

Dataset Class Number	Blood 8	Breast 2	Derma 7	Organ 11	Pathology 9	Pneumonia 2	Retina 5
ViT [2]	0.985	0.859	0.883	0.961	0.965	0.877	0.642
LoRA [3]	0.988	0.885	0.859	0.968	0.953	0.874	0.620
SVD-LoRA	0.991	0.878	0.866	0.968	0.945	0.872	0.655
PMC-CLIP [7]	0.070	0.660	0.374	0.119	0.162	0.468	0.088
ViT Averaging [15]	0.286	0.744	0.196	0.451	0.580	<u>0.846</u>	0.220
LoRA Averaging	0.315	0.782	0.602	0.353	0.218	0.633	0.510
Task Arithmetic [4]	0.078	0.718	0.201	0.967	0.342	0.402	0.405
Ties-Merging [16]	0.215	0.731	0.129	<u>0.737</u>	0.669	0.782	0.147
PEM Composition [19]	<u>0.385</u>	<u>0.814</u>	<u>0.685</u>	0.327	0.296	0.785	<u>0.530</u>
MagMax [8]	0.137	0.731	0.052	0.117	0.101	0.764	0.420
Med-LEGO (ours)	0.814	0.872	0.739	0.567	<u>0.606</u>	0.849	0.593

3.4 Performance of in-domain Tasks Merging

We test in-domain information sharing performance using three different chest X-ray datasets, with results shown in Table 2. The NIH-CXR14 dataset, being a multi-label task, shows lower accuracy across all methods because it is more challenging compared to other single-label tasks. Our proposed Med-LEGO method effectively balances information across the three chest X-ray datasets, while enhancing performance on more challenging tasks. At the same time, it maintains nearly the same performance on simpler tasks with minimal loss. In contrast, other methods are misled by the more difficult NIH-CXR14 task, leading to catastrophic forgetting on the other two datasets. Experimental results demonstrate that Med-LEGO efficiently performs information fusion in in-domain medical image diagnosis tasks.

3.5 Generalization Performance on New Tasks

To verify whether the new weights obtained from merging using Med-LEGO have generalization capability, we introduce three additional new datasets from different tasks for fine-tuning. This tests whether the model weights transition from natural image pre-training to medical image pre-training. We use the fine-tuning strategy mentioned in Section 3.2, and we include the best test accuracy of the ViT fine-tuned individually on each dataset as a reference for natural image pre-training weights.

We compare the ViT-based merging method, ViT Averaging, and the LoRA-based merging method, PEM Composition. What’s more, Med-LEGO (General) comes from Section 3.3, and Med-LEGO (CXR) comes from Section 3.4. The

Table 2. Comparison of the accuracy of model merging methods across 3 in-domain chest X-ray datasets, where gray section represents fine-tuning for each task individually. The remaining section represents testing a single merged model on all tasks. **Bold** indicates the best result, and the underlined number denotes the second-best result.

Dataset	Pneumonia		NIH-CXR14		Tuberculosis	
Class Number	2		14		2	
Metrics	ACC	AUC	ACC	AUC	ACC	AUC
ViT [2]	0.877	0.950	0.361	0.737	0.812	0.894
LoRA [3]	0.874	0.972	0.368	0.783	0.818	0.916
SVD-LoRA	0.872	0.967	0.369	0.780	0.818	0.915
ViT Averaging [15]	0.354	0.238	0.385	0.710	0.515	0.542
Task Arithmetic [4]	0.372	0.239	0.385	<u>0.703</u>	0.530	0.527
Ties-Merging [16]	0.330	0.197	0.385	0.686	0.500	0.472
PEM Composition [19]	<u>0.633</u>	<u>0.887</u>	0.385	0.579	<u>0.773</u>	<u>0.806</u>
MagMax [8]	0.579	0.514	0.001	0.493	0.712	0.783
Med-LEGO (ours)	0.845	0.966	<u>0.382</u>	0.710	0.803	0.897

experimental results are shown in Fig 2, we see that the weights obtained from Med-LEGO show a more stable training process and better convergence on the new medical image datasets. This further demonstrates that the model weights merged by Med-LEGO represent general medical image information, transforming the original natural image pre-training model into a pre-trained model better suited for medical image diagnosis.

Notably, the general weight obtained from chest X-ray datasets performs better on Tuberculosis than the general weight obtained from cross-domain data. This demonstrates that Med-LEGO is equally effective in enhancing specialist models for similar tasks. Whats more, compared to the LoRA-based merging method, PEM Composition, Med-LEGO retains the low-rank adaptation structure after merging, which leads to lower training resource consumption during fine-tuning on subsequent new tasks.




Fig. 2. Performance on fine-tuning with a new dataset. We used the merged model weights from Section 3.3 and Section 3.4 to fine-tune for the new task and present the accuracy results on the test set as the training epochs increase.

4 Discussion and Conclusion

Model editing without training is essential for advancing general models in medical imaging. Given the privacy concerns and limited availability of medical image datasets, leveraging model weights fine-tuned on specific tasks to merge information offers a convenient solution. This approach enables continuous updates to the original pre-trained model, facilitating the creation of a universal medical foundation model. Our proposed Med-LEGO method enhances task adaptation efficiency through SVD-LoRA and represents model capabilities with less than 1% of the original models parameters, significantly reducing file size. Additionally, our SVD merge method effectively integrates cross-domain and in-domain information from different datasets, resolving conflicts between them. As a result, the model updated through Med-LEGO exhibits stronger adaptability to new medical imaging tasks.

Acknowledgments. This work was partially supported by National Natural Science Foundation of China (62131015), AI4S Initiative and HPC Platform of ShanghaiTech University.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
2. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
3. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
4. Ilharco, G., Ribeiro, M.T., Wortsman, M., Gururangan, S., Schmidt, L., Hājishirzi, H., Farhadi, A.: Editing models with task arithmetic. arXiv preprint arXiv:2212.04089 (2022)
5. Jin, X., Ren, X., Preotiuc-Pietro, D., Cheng, P.: Dataless knowledge fusion by merging weights of language models. arXiv preprint arXiv:2212.09849 (2022)
6. Kellgren, J., Lawrence, J.: Radiological assessment of osteo-arthritis. *Annals of the rheumatic diseases* **16**(4), 494 (1957)
7. Lin, W., Zhao, Z., Zhang, X., Wu, C., Zhang, Y., Wang, Y., Xie, W.: Pmc-clip: Contrastive language-image pre-training using biomedical documents. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 525–536. Springer (2023)
8. Marczak, D., Twardowski, B., Trzcíński, T., Cygert, S.: Magmax: Leveraging model merging for seamless continual learning. In: European Conference on Computer Vision. pp. 379–395. Springer (2024)

9. Matena, M.S., Raffel, C.A.: Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems* **35**, 17703–17716 (2022)
10. Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J., Rajpurkar, P.: Foundation models for generalist medical artificial intelligence. *Nature* **616**(7956), 259–265 (2023)
11. Razzak, M.I., Naz, S., Zaib, A.: Deep learning for medical image processing: Overview, challenges and the future. *Classification in BioApps: Automation of decision making* pp. 323–350 (2017)
12. Steiner, A., Kolesnikov, A., , Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L.: How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270* (2021)
13. Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., Hurst, R.T., Kendall, C.B., Gotway, M.B., Liang, J.: Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE transactions on medical imaging* **35**(5), 1299–1312 (2016)
14. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2097–2106 (2017)
15. Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al.: Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: *International conference on machine learning*. pp. 23965–23998. PMLR (2022)
16. Yadav, P., Tam, D., Choshen, L., Raffel, C.A., Bansal, M.: Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems* **36**, 7093–7115 (2023)
17. Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., Ni, B.: Medmnist v2- a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Scientific Data* **10**(1), 41 (2023)
18. Yu, L., Yu, B., Yu, H., Huang, F., Li, Y.: Language models are super mario: Absorbing abilities from homologous models as a free lunch. In: *Forty-first International Conference on Machine Learning* (2024)
19. Zhang, J., Liu, J., He, J., et al.: Composing parameter-efficient modules with arithmetic operation. *Advances in Neural Information Processing Systems* **36**, 12589–12610 (2023)
20. Zhang, Q., Chen, M., Bukharin, A., Karampatziakis, N., He, P., Cheng, Y., Chen, W., Zhao, T.: Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512* (2023)
21. Zhu, Y., Shen, Z., Zhao, Z., Wang, S., Wang, X., Zhao, X., Shen, D., Wang, Q.: Melo: Low-rank adaptation is better than fine-tuning for medical image diagnosis. In: *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5. IEEE (2024)