
Beyond the Permutation Symmetry of Transformers: The Role of Rotation for Model Fusion

Binchi Zhang^{* 1} Zaiyi Zheng^{* 1} Zhengzhang Chen² Jundong Li¹

Abstract

Symmetry in the parameter space of deep neural networks (DNNs) has proven beneficial for various deep learning applications. A well-known example is the permutation symmetry in Multi-Layer Perceptrons (MLPs), where permuting the rows of weight matrices in one layer and applying the inverse permutation to adjacent layers yields a functionally equivalent model. While permutation symmetry fully characterizes the equivalence set for MLPs, its discrete nature limits its utility for transformers. In this paper, we introduce rotation symmetry, a novel form of parameter space symmetry for transformers that generalizes permutation symmetry by rotating parameter matrices in self-attention layers. Unlike permutation symmetry, rotation symmetry operates in a continuous domain, thereby significantly expanding the equivalence set for transformers. Based on this property, we propose a theoretically optimal parameter matching algorithm as a plug-and-play module to enhance model fusion. We evaluate our approach using pre-trained transformers across diverse natural language and vision tasks. Experimental results demonstrate that our rotation symmetry-based matching algorithm substantially improves model fusion, highlighting the potential of parameter space symmetry to facilitate model fusion. Our code is available on <https://github.com/zhengzaiyi/RotationSymmetry>.

1. Introduction

Parameter space symmetry is an intriguing property of neural networks that has garnered increasing attention in recent years (Du et al., 2018; Armenta & Jodoin, 2021; Kunin et al.,

2021; Simsek et al., 2021; Entezari et al., 2022; Grigsby et al., 2023; Lim et al., 2024b; Ziyin et al., 2024; 2025). One of the most studied forms of parameter space symmetry is permutation symmetry (Ainsworth et al., 2023; Entezari et al., 2022). For instance, in a two-layer MLP, permuting the rows of the weight matrix in the first layer and applying the corresponding inverse permutation to the second layer results in a functionally equivalent model, *i.e.*, the outputs of the original and permuted models remain identical for any given input (Ainsworth et al., 2023). All functionally equivalent models corresponding to weight permutations form an equivalence set, which provides theoretical insights into neural network optimization, such as the linear mode connectivity of loss landscapes (Entezari et al., 2022; Zhou et al., 2023c; Ferbach et al., 2024). In addition, permutation symmetry has also proven helpful in advancing neural network applications, such as model fusion (Singh & Jaggi, 2020; Ainsworth et al., 2023) and optimization (Zhao et al., 2024; Zamir et al., 2025).

Although parameter space symmetry has been extensively studied in classical neural network architectures, such as MLPs and CNNs, the understanding of its application in transformers (Vaswani et al., 2017) remains limited. Transformers have seen rapid advancements in recent years, achieving remarkable success in a wide range of applications (Yun et al., 2019; Lewis et al., 2020; Raffel et al., 2020; Clark et al., 2020; Zhou et al., 2021; He et al., 2024; Zhu et al., 2024; Zheng et al., 2024; Zhang et al., 2025; Wei & Liu, 2025). The transformer architecture is built upon two primary submodules: feedforward networks and self-attention layers. The feedforward network, which is structurally similar to MLPs, naturally inherits the permutation symmetry that has been extensively studied in the existing literature. Self-attention layers, on the other hand, involve a unique attention mechanism powered by matrix products of queries, keys, and values, which introduce additional potentials for symmetry beyond permutations. Permutation symmetries limit the equivalence set of neural networks to *discrete operations*, which aligns well with MLPs due to their element-wise activations (*e.g.*, ReLU (Glorot et al., 2011)). In contrast, the *continuous nature* of the matrix operations in self-attention layers necessitates more flexible operations to fully characterize their equivalence set.

^{*}Equal contribution ¹University of Virginia ²NEC Laboratories America. Correspondence to: Jundong Li <jundong@virginia.edu>.

In this paper, we introduce **rotation symmetry**, a novel form of parameter space symmetry for self-attention layers in transformers. Specifically, we analyze the query and key matrices jointly and demonstrate that applying a rotation to the query matrix, followed by the corresponding inverse rotation to the key matrix, preserves the query-key product. Additionally, we find that the same rotation rule can be applied to the value and output matrices. Our findings provide novel insights into the functional invariance of the attention mechanism and extend the permutation symmetry (Singh & Jaggi, 2020; Wang et al., 2020; Tatro et al., 2020; Entezari et al., 2022; Ainsworth et al., 2023; Imfeld et al., 2024) from discrete spaces to continuous spaces, which significantly extends the scope of parameter space symmetries for transformers.

To further demonstrate the benefits of rotation symmetry for transformers, we explore its utility in **model fusion**. The goal of model fusion is to merge multiple well-trained end models in the parameter space to produce a single merged model with improved overall utility. Model fusion is widely adopted across various settings, such as hyperparameter tuning (where end models are trained on the same benchmark) and multi-task learning (where end models are trained with different tasks) (Jin et al., 2023). Unlike ensemble learning (Dietterich, 2000; Lakshminarayanan et al., 2017; Sagi & Rokach, 2018; Dong et al., 2020), model fusion can work in a data-agnostic manner, making it suitable for privacy-sensitive scenarios such as federated learning (Yurochkin et al., 2019; Wang et al., 2020).

The existing literature has demonstrated that the performance of model fusion is closely tied to the distance between the end models (Wortsman et al., 2022). Inspired by this finding, we propose a parameter matching algorithm that selects functionally equivalent end models from the equivalence class determined by rotation symmetry. This approach ensures that the selected representative models are closer in parameter space, resulting in a smaller inner distance. To achieve this, we formulate the problem of parameter matching as an optimization problem with orthogonal constraints. Leveraging the continuous nature of rotation symmetry, we propose a closed-form solution to this problem. Our parameter matching algorithm is highly efficient, easy to implement, and can be seamlessly incorporated as a plug-and-play module for model fusion. To evaluate its effectiveness, we conduct extensive experiments with pre-trained transformers on real-world NLP and vision tasks. The experimental results demonstrate that incorporating rotation symmetry into parameter matching improves model fusion effectively and efficiently. Furthermore, additional experiments reveal that even matching a subset of parameters can lead to notable performance improvements, highlighting the practical utility of our approach. Our contributions are threefold:

- We introduce a novel rotation symmetry for the attention mechanism in transformers, extending the concept of symmetry to a continuous space.
- Building on rotation symmetry, we propose a theoretically optimal parameter matching algorithm that improves the effectiveness of model fusion in transformers.
- Through extensive experiments, we validate the efficacy of our proposed parameter matching algorithm, demonstrating its potential to advance model fusion through parameter space symmetry.

2. Preliminaries

To facilitate further discussion, we begin by defining the notations. Let $\mathbf{A}[i, :]$, $\mathbf{A}[:, j]$, and $\mathbf{A}[i, j]$ denote the i -th row vector, the j -th column vector, and the (i, j) -th element of the matrix \mathbf{A} , respectively. Additionally, let $\mathbf{1}$ denote an all-ones vector. We now illustrate the concept of permutation symmetry in neural networks, using an MLP as an example. Consider an L -layer MLP model defined as:

$$f_{\mathbf{W}}(\mathbf{X}) = \mathbf{Z}^{(L)}, \quad \mathbf{Z}^{(l)} = \sigma(\mathbf{Z}^{(l-1)}(\mathbf{W}^{(l)})^\top + \mathbf{b}^{(l)}), \quad (1)$$

where $\mathbf{Z}^{(0)} = \mathbf{X}$ is the input feature matrix, $\mathbf{b}^{(l)}$ is the bias vector ($\mathbf{b}^{(l)}$ is a row vector and can be seen as broadcast over all rows in Equation (1)), $\mathbf{W} = \{\mathbf{W}^{(l)}, \mathbf{b}^{(l)}\}_{l=1, \dots, L}$ collects all learnable parameters, and $\sigma(\cdot)$ stands for a non-linear activation function, such as ReLU (Nair & Hinton, 2010). A loss function $\mathcal{L}(f_{\mathbf{W}}(\mathbf{X}), \mathbf{Y})$ is used to measure the distance between the model prediction and the ground truth label \mathbf{Y} .

To analyze the permutation symmetry in the MLP model, let $\mathbf{P} \in \mathcal{P}$ be a permutation matrix, where $\mathbf{P}[i, j] \in \{0, 1\}$ and $\mathbf{P}[i, :]\mathbf{1} = \mathbf{P}[:, j]^\top \mathbf{1} = 1$ for any i, j . All permutation matrices are orthogonal (Strang, 1976), satisfying $\mathbf{P}^\top = \mathbf{P}^{-1}$. Consequently, for layer l and $l + 1$, we have:

$$\begin{aligned} \mathbf{Z}^{(l+1)} &= \sigma(\sigma(\mathbf{Z}^{(l-1)}(\mathbf{W}^{(l)})^\top + \mathbf{b}^{(l)})(\mathbf{W}^{(l+1)})^\top + \mathbf{b}^{(l+1)}) \\ &= \sigma(\sigma(\mathbf{Z}^{(l-1)}(\mathbf{P}^\top \mathbf{W}^{(l)})^\top + \mathbf{b}^{(l)} \mathbf{P})(\mathbf{W}^{(l+1)} \mathbf{P})^\top + \mathbf{b}^{(l+1)}). \end{aligned} \quad (2)$$

The third equal sign holds because the element-wise activation function σ is decoupled from column permutation (*i.e.*, being multiplied by \mathbf{P}). Based on Equation (2), it follows that for layers l and $l + 1$, the mappings $\mathbf{W}^{(l)} \rightarrow \mathbf{P}^\top \mathbf{W}^{(l)}$, $\mathbf{b}^{(l)} \rightarrow \mathbf{b}^{(l)} \mathbf{P}$, $\mathbf{W}^{(l+1)} \rightarrow \mathbf{W}^{(l+1)} \mathbf{P}$ preserves the output $\mathbf{Z}^{(l+1)}$ for any input $\mathbf{Z}^{(l-1)}$. For each pair of adjacent layers, a similar mapping exists independently based on a specific permutation matrix, denoted as $\mathbf{P}^{(l)}$ for layers l and $l + 1$. Consequently, for the entire L -layer MLP, there exists a mapping that preserves the model's predictions for any input \mathbf{X} : $\mathbf{W}^{(l)} \rightarrow (\mathbf{P}^{(l)})^\top \mathbf{W}^{(l)} \mathbf{P}^{(l-1)}$, $\mathbf{b}^{(l)} \rightarrow \mathbf{b}^{(l)} \mathbf{P}^{(l)}$. Equivalently, if we define $\mathbf{W}' = \{(\mathbf{P}^{(l)})^\top \mathbf{W}^{(l)} \mathbf{P}^{(l-1)}, \mathbf{b}^{(l)} \mathbf{P}^{(l)}\}_{l=1, \dots, L}$ for any $\mathbf{P}^{(l)} \in \mathcal{P}$

($l = 1, \dots, L - 1$) and $P^{(0)} = P^{(L)} = I$, then we have $f_{W'}(X) = f_W(X)$ for any X . This phenomenon is referred to as the *permutation symmetry* of the parameter space (Godfrey et al., 2022; Hecht-Nielsen, 1990; Navon et al., 2023; Rossi et al., 2023; Simsek et al., 2021). Leveraging permutation symmetry allows us to identify an equivalence class of functionally equivalent model parameters, which is known as *permutation invariance* (Ainsworth et al., 2023; Entezari et al., 2022; Lubana et al., 2023). We denote the equivalence relation induced by permutation invariance as π , where, for example, $W' = \pi(W)$ represents the equivalence between the original parameters W and the permuted parameters W' .

3. Parameter Space Symmetry of Transformers

Transformers (Vaswani et al., 2017) have revolutionized deep learning with their ability to handle sequential data effectively, particularly in natural language processing (NLP) and other fields (Brown et al., 2020; Devlin et al., 2019; Liu et al., 2021; Radford et al., 2021). Their success is primarily driven by two key components¹: feedforward networks and self-attention layers (Vaswani et al., 2017). To better understand the parameter space symmetry of transformers, we examine these two core modules individually in the following sections.

3.1. Permutation Symmetry of Feedforward Networks

We first look at the feedforward networks. The feedforward network adopted in the transformer blocks is a two-layer MLP model which can be written as

$$FFN(X) = LN(ReLU(XW_i^\top + b_i)W_o^\top + b_o + X), \quad (3)$$

where LN denotes the Layer Normalization operator (Ba et al., 2016). Different from (Vaswani et al., 2017), we include the residual connection (He et al., 2016) and layer normalization modules into the formula of the feedforward network (and the self-attention layer mentioned later). According to Equation (2) and the analysis in Preliminaries, the feedforward networks have the permutation symmetry property and the equivalence class determined by permutation invariance is defined as

$$W_i \rightarrow P^\top W_i, b_i \rightarrow b_i P, W_o \rightarrow W_o P, b_o \rightarrow b_o, \quad (4)$$

where $P \in \mathcal{P}$ is a permutation matrix. It is worth noting that the permutations of different feedforward networks in a transformer are independent due to the scalable modular design. Consequently, we are able to flexibly compute the permutation invariance equivalence class of each FFN module in a transformer model.

¹We follow the architecture in the original transformer paper (Vaswani et al., 2017).

3.2. Rotation Symmetry of Self-attention Layers

We then focus on the self-attention layers. In this paper, we introduce the *rotation symmetry* of self-attention layers. For MLPs, we have to switch the order of multiplying P and passing the activation function σ (the third equal sign in Equation (2)), requiring the matrix P to be a permutation matrix. In contrast, the self-attention layer does not contain an element-wise activation function, enabling a wider range of the matrix P in self-attention layers. The self-attention layer can be written as

$$ATTN(X) = LN \left(Cat_{h=1}^H \left\{ X_{QKV}^h \right\} W_O^\top + b_O + X \right), \\ X_{QKV}^h = Sftmx \left(X_Q^h (X_K^h)^\top / \sqrt{d_k} \right) \cdot X_V^h,$$

where Cat stands for the operator concatenating the outputs of multi-head attention, $Sftmx(\cdot)$ denotes the softmax operator, H denotes the number of multi-heads, and the subscripts Q , K , V , and O denote Query, Key, Value, and Output, respectively.

We first transform the query and key matrices as

$$\begin{aligned} X_Q^h (X_K^h)^\top &= (X(W_Q^h)^\top + b_Q^h)(X(W_K^h)^\top + b_K^h)^\top \\ &= (X(W_Q^h)^\top + b_Q^h)RR^\top (X(W_K^h)^\top + b_K^h)^\top \\ &= (X(R^\top W_Q^h)^\top + b_Q^h R)(X(R^\top W_K^h)^\top + b_K^h R)^\top, \end{aligned}$$

where R is a rotation matrix, *i.e.*, $RR^\top = I$. It is worth noting that each multi-head corresponds to a specific rotation matrix R . Let $W_O = [W_O^1 \ W_O^2 \ \dots \ W_O^H]$ and we then rewrite the concatenating operation (with the product by W_O) as $\sum_{h=1}^H S(\cdot \cdot) (X(W_V^h)^\top + b_V^h)(W_O^h)^\top$. Similar to the Q-K case, we can transform the value and output matrices as $(X(W_V^h)^\top + b_V^h)(W_O^h)^\top = (X(W_V^h)^\top + b_V^h)RR^\top (W_O^h)^\top = (X(R^\top W_V^h)^\top + b_V^h R)(W_O^h R)^\top$ for each multi-head h , where R is a rotation matrix. Finally, we derive an equivalence class of the parameters in a self-attention layer determined by rotation invariance as

$$\begin{aligned} W_Q^h &\rightarrow (R_{qk}^h)^\top W_Q^h, b_Q^h \rightarrow b_Q^h R_{qk}^h, \\ W_K^h &\rightarrow (R_{qk}^h)^\top W_K^h, b_K^h \rightarrow b_K^h R_{qk}^h, W_V^h \rightarrow (R_{vo}^h)^\top W_V^h, \\ b_V^h &\rightarrow b_V^h R_{vo}^h, W_O^h \rightarrow W_O^h R_{vo}^h, b_O \rightarrow b_O, \end{aligned} \quad (5)$$

where R_{qk}^h and R_{vo}^h are rotation matrices for $h = 1, \dots, H$. The progress from permutation to rotation extends the symmetry of transformers to a continuous space and enhances our understanding of the parameter space symmetry of attention mechanism. The denseness of the symmetry allows for the choice of better invariant models to analyze transformers' loss landscapes. For better clarity, we provide an intuitive illustration of the rotation symmetry in Figure 1.

4. Symmetry for Model Fusion

In this section, we explore the benefit of the rotation symmetry of transformers in model fusion (Li et al., 2023; Matena

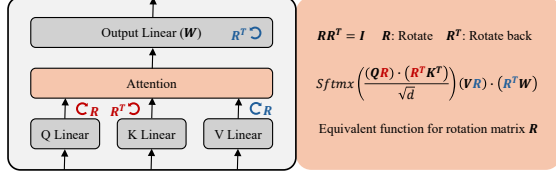


Figure 1. The rotation symmetry of self-attention layers.

& Raffel, 2022; Wortsman et al., 2022; Yadav et al., 2023; Jin et al., 2023; Daheim et al., 2024; Yang et al., 2024). Model fusion is proposed to merge multiple given end models trained in different settings (e.g., upon different datasets and hyperparameter settings) in the parameter space to improve model utility and robustness. Compared with ensemble learning (Dietterich, 2000; Sagi & Rokach, 2018; Lakshminarayanan et al., 2017), model fusion has a lower inference-stage complexity without requiring access to the training data. Most existing methods of model fusion conduct a weighted averaging of different end models, e.g., direct averaging (Wortsman et al., 2022), Fisher-weighted averaging (Matena & Raffel, 2022), and regression-mean averaging (Jin et al., 2023). We next show the potential of exploiting the permutation and rotation symmetry as a plug-and-play module to improve the model fusion techniques.

4.1. Background and Motivation

Let W_1, \dots, W_k denote k different end models (after training or pre-training) with the same architecture. The goal of model fusion is to merge the given k end models in the parameter space and obtain a single model. If the given models are trained over different datasets, we can expect the merged model to have better utility and out-of-distribution robustness (Jin et al., 2023). The theoretical results in previous literature (Wortsman et al., 2022) have shown that *strong convexity* and *closer end models* can boost the utility of direct model fusion. Consequently, the primary advantage of permutation (and rotation) symmetry applying to model fusion is that we can substitute the end models with corresponding carefully chosen equivalent models (in the equivalence class determined by permutation and rotation invariance) to make the selected end models more concentrated, i.e., closer to each other. This step is usually called **parameter matching** (Singh & Jaggi, 2020; Wang et al., 2020; Ainsworth et al., 2023), aka parameter or neuron alignment. Parameter matching, which aims to reduce the distance between end models, can naturally yield closer end models and improve model fusion performance. On the other hand, different end models can lie in the basins of different local optimums regarding the highly non-convex nature of transformers. Previous studies have verified that parameter matching can merge different end models toward a single low-loss basin, i.e., the loss value along the linear interpolation between matched models shows an approxi-

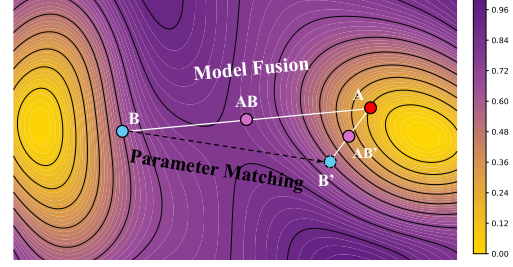


Figure 2. An intuitive example of the usage of parameter space symmetry for model fusion. The background shows the contour map of the loss landscape in the model parameter space. A and B are the original end models to be merged, AB is the result of naive model fusion, and AB' is the result of model fusion with parameter matching.

mately flat or convex curve (Entezari et al., 2022; Ainsworth et al., 2023). This property is called *linear mode connectivity* and is regarded as a weak form of convexity (Ainsworth et al., 2023). As a result, parameter matching can also help improve the convexity of the objective in the area adjacent to the end models. We showcase an intuitive example of using parameter space symmetry to improve model fusion in Figure 2 to illustrate the intuition behind parameter matching. We can observe that the merged model after parameter matching (AB') has a lower loss value, i.e., better utility than the naive merged model (AB).

4.2. Parameter Matching

Next, our primary goal is to develop a practical parameter matching algorithm to minimize the distance between different end models. We begin by merging two models W_1 and W_2 .

Matching two FFNs. Let $\{W_{i_k}, b_{i_k}, W_{o_k}, b_{o_k}\}_{k=1,2}$ denote the model parameters in the two feedforward networks to be merged where k denote the index of the networks. We have already derived the equivalence relation of parameters in the feedforward networks determined by permutation invariance as Equation (4). Based on Equation (4), we can formulate parameter matching as an optimization problem as follows.

$$\begin{aligned} \arg\min_{P_1, P_2 \in \mathcal{S}} & \|P_1^\top W_{i_1} - P_2^\top W_{i_2}\|_F^2 \\ & + \|b_{i_1} P_1 - b_{i_2} P_2\|_F^2 + \|W_{o_1} P_1 - W_{o_2} P_2\|_F^2. \end{aligned} \quad (6)$$

As shown in Equation (6), the goal is to find the functionally equivalent models from the equivalence classes $\pi(W_1)$ and $\pi(W_2)$ determined by P_1 and P_2 , which has the smallest ℓ_2 distance in the parameter space. Following the method (Ainsworth et al., 2023), we reformulate the optimization problem shown in Equation (6) as a *linear assignment problem*:

$$\arg\max_{P \in \mathcal{S}} \langle P, W_{i_1} W_{i_2}^\top + b_{i_1}^\top b_{i_2} + W_{o_1}^\top W_{o_2} \rangle_F. \quad (7)$$

Linear assignment problems such as Equation (7) have been well-studied in previous literature (Martello & Toth, 1987; Burkard & Cela, 1999) and can be solved precisely by Hungarian Algorithm (Martello & Toth, 1987). After calculating the value of P , we substitute $P = P_1 P_2^{-1}$ back to Equation (4) and let $P_2 = I$ (for simplicity), then we obtain the matched parameter as $W_{i_1} \rightarrow P^\top W_{i_1}$, $b_{i_1} \rightarrow b_{i_1} P$, $W_{O_1} \rightarrow W_{O_1} P$, $b_{O_1} \rightarrow b_{O_1}$. Parameters of the second model $\{W_{i_2}, b_{i_2}, W_{O_2}, b_{O_2}\}$ remain unchanged, which makes the second model an anchor model for matching.

Matching two ATTNs. Let the parameters in the two self-attention layers to be merged be $\{W_{Q_k}^i, b_{Q_k}^i, W_{K_k}^i, b_{K_k}^i, W_{V_k}^i, b_{V_k}^i, W_{O_k}^i, b_{O_k}^i\}_{k=1,2;i=1\dots H}$ where k denote the index of the layers. The equivalence relation of these parameters in the self-attention layers in terms of rotation invariance is shown in Equation (5). Similar to matching FFNs, our goal is to match the parameters of one (source) attention layer to the other (target) attention layer where the matched parameters are supposed to be *closest* to the target parameters while being functionally equivalent when feeding with any input data, *i.e.*, in the equivalence class determined by rotation invariance. The goal can be formulated as an optimization problem.

$$\begin{aligned} \min_{R_{qk}^h, R_{vo}^h \in \mathcal{R}} & \left\| \begin{bmatrix} (W_{Q_1}^h)^\top & (W_{Q_2}^h)^\top \\ b_{Q_1}^h & b_{Q_2}^h \end{bmatrix} \begin{bmatrix} R_{qk1}^h \\ -R_{qk2}^h \end{bmatrix} \right\|_F^2 \\ & + \left\| \begin{bmatrix} (W_{K_1}^h)^\top & (W_{K_2}^h)^\top \\ b_{K_1}^h & b_{K_2}^h \end{bmatrix} \begin{bmatrix} R_{qk1}^h \\ -R_{qk2}^h \end{bmatrix} \right\|_F^2 \\ & + \left\| \begin{bmatrix} (W_{V_1}^h)^\top & (W_{V_2}^h)^\top \\ b_{V_1}^h & b_{V_2}^h \end{bmatrix} \begin{bmatrix} R_{vo1}^h \\ -R_{vo2}^h \end{bmatrix} \right\|_F^2 \\ & + \left\| \begin{bmatrix} W_{O_1}^h & W_{O_2}^h \\ 0 & 0 \end{bmatrix} \begin{bmatrix} R_{vo1}^h \\ -R_{vo2}^h \end{bmatrix} \right\|_F^2, \end{aligned} \quad (8)$$

where \mathcal{R} denotes the set of rotation matrices. To solve this problem, we divide Equation (8) into two separate optimization problems (the first line in terms of R_{qk1}^h, R_{qk2}^h as the first objective and the second line in terms of R_{vo1}^h, R_{vo2}^h as the second objective). Considering the similar formulation of these two optimization problems, in this paper, we propose the following theorem to solve both problems.

Theorem 4.1. *The following optimization problem has a closed-form solution.*

$$\begin{aligned} \min_{R_1, R_2 \in \mathcal{R}} & \left\| \begin{bmatrix} W_{Q_1}^\top & W_{Q_2}^\top \\ b_{Q_1} & b_{Q_2} \end{bmatrix} \begin{bmatrix} R_1 \\ -R_2 \end{bmatrix} \right\|_F^2 \\ & + \left\| \begin{bmatrix} W_{K_1}^\top & W_{K_2}^\top \\ b_{K_1} & b_{K_2} \end{bmatrix} \begin{bmatrix} R_1 \\ -R_2 \end{bmatrix} \right\|_F^2. \end{aligned} \quad (9)$$

The solution is given by

$$R_1 = UV^\top, R_2 = I, \quad (10)$$

Algorithm 1 Matching two self-attention layers.

Input: Model parameters of the attention layers $\{W_{Q_k}^h, b_{Q_k}^h, W_{K_k}^h, b_{K_k}^h, W_{V_k}^h, b_{V_k}^h, W_{O_k}^h, b_{O_k}^h\}_{k=1,2;h=1,\dots,H}$.
Output: The optimally matched parameters from source parameters ($k = 1$) to anchor parameters ($k = 2$, maintain unchanged).

- 1: QK solution: $R_{qk1}^h = U_{qk}^h (V_{qk}^h)^\top$, $R_{qk2}^h = I$, where $U_{qk}^h \Sigma_{qk}^h (V_{qk}^h)^\top = W_{Q_1}^h (W_{Q_2}^h)^\top + W_{K_1}^h (W_{K_2}^h)^\top + (b_{Q_1}^h)^\top b_{Q_2}^h + (b_{K_1}^h)^\top b_{K_2}^h$.
- 2: QK matching: $W_{Q_1}^h \rightarrow (R_{qk1}^h)^\top W_{Q_1}^h$, $W_{K_1}^h \rightarrow (R_{qk1}^h)^\top W_{K_1}^h$, $b_{Q_1}^h \rightarrow b_{Q_1}^h R_{qk1}^h$, $b_{K_1}^h \rightarrow b_{K_1}^h R_{qk1}^h$.
- 3: VO solution: $R_{vo1}^h = U_{vo}^h (V_{vo}^h)^\top$, $R_{vo2}^h = I$, where $U_{vo}^h \Sigma_{vo}^h (V_{vo}^h)^\top = W_{V_1}^h (W_{V_2}^h)^\top + (W_{O_1}^h)^\top W_{O_2}^h + (b_{V_1}^h)^\top b_{V_2}^h$.
- 4: VO matching: $W_{V_1}^h \rightarrow (R_{vo1}^h)^\top W_{V_1}^h$, $W_{O_1}^h \rightarrow W_{O_1}^h R_{vo1}^h$, $b_{V_1}^h \rightarrow b_{V_1}^h R_{vo1}^h$.

where I is the identity matrix and $U \Sigma V^\top = W_{Q_1} W_{Q_2}^\top + W_{K_1} W_{K_2}^\top + b_{Q_1}^\top b_{Q_2} + b_{K_1}^\top b_{K_2}$ is the result of eigendecomposition.

We leave a detailed proof of Theorem 4.1 in Appendix A. According to Theorem 4.1, we can obtain the algorithm to match two self-attention layers shown in Algorithm 1. Algorithm 1 can be seen as an adaptation of the Kabsch algorithm (Kabsch, 1976; Umeyama, 1991). Without loss of generality, we let the parameters in the second ($k = 2$) self-attention layer be the anchor and conduct rotation for the first layer ($k = 1$). The denseness of rotation symmetry helps reduce the distance between the end models after parameter matching as Algorithm 1 shows.

Rescaling Matching. We find that the rescaling symmetry (Neyshabur et al., 2015; Meng et al., 2019; Godfrey et al., 2022; Kalogeropoulos et al., 2024) can be integrated with our proposed rotation symmetry in the parameter matching algorithm. For instance, for the (simplified) Q-K product $W_Q W_K$, we can find a rescaling operation that preserves the functionality of the Q-K product $a W_Q \cdot \frac{1}{a} W_K$ where $a \neq 0$ is a real number. By adding a scalar variable to each parameter matrix, we can formulate the objective of rescaling symmetry as follows (taking the Q-K product as an example).

$$\begin{aligned} \min_a & \|a W_{Q_1}' - W_{Q_2}'\|^2 + \|a b_{Q_1}' - b_{Q_2}'\|^2 \\ & + \|W_{K_1}'/a - W_{K_2}'\|^2 + \|b_{K_1}'/a - b_{K_2}'\|^2, \end{aligned} \quad (11)$$

where a is the rescaling variable and $'$ denotes the parameters after rotation symmetry-based matching. Here, we still set model 2 as the anchor model and conduct the rescaling operation on model 1 to align with model 2. We can easily

Table 1. Experimental results of in-domain (Emotion and NER) and out-of-domain (NER-CoNLL) model fusion for two base language models over Emotion and NER tasks.

	Emotion		NER		NER-CoNLL	
	Roberta	Deberta	Roberta	Deberta	Roberta	Deberta
Simple	35.87	2.99	60.88	27.54	26.86	10.80
+match	35.87	2.99	60.88	31.31	26.87	23.30
Fisher	44.02	35.95	54.55	33.20	23.06	12.53
+match	44.05	35.98	54.58	33.83	23.05	18.33
Regmean	35.87	2.99	60.88	27.54	26.86	10.80
+match	39.95	2.99	60.88	31.31	26.87	14.06

solve the optimality condition of Equation (11) as

$$\begin{aligned}
 & \left(\|W'_{Q_1}\|^2 + \|b'_{Q_1}\|^2 \right) a^4 - \|W'_{K_1}\|^2 - \|b'_{K_1}\|^2 \\
 & - (\langle W'_{Q_1}, W'_{Q_2} \rangle + \langle b'_{Q_1}, b'_{Q_2} \rangle) a^3 \\
 & + (\langle W'_{K_1}, W'_{K_2} \rangle + \langle b'_{K_1}, b'_{K_2} \rangle) a = 0.
 \end{aligned} \tag{12}$$

The roots of Equation (12) can be derived using numerical methods as the value of a . We then conduct the rescaling operation to model 1 as $W'_{Q_1} \rightarrow aW'_{Q_1}$, $b'_{Q_1} \rightarrow ab'_{Q_1}$, $W'_{K_1} \rightarrow \frac{1}{a}W'_{K_1}$, $b'_{K_1} \rightarrow \frac{1}{a}b'_{K_1}$. It is worth noting that the rescaling symmetry-based matching is conducted after the rotation symmetry-based matching. Using the rescaling operation, we extend the rotation matrices to orthogonal matrices with larger norms. Nevertheless, the end models are close to each other in practical scenarios so the value of a is usually close to 1.

Complexity. We next provide a brief analysis of the complexity of our proposed parameter matching algorithm. We let the hidden dimension of the target transformer be d , the layer of the target transformer be L , and the number of attention heads be H . For the parameter matching of feedforward networks, the linear assignment problem can be solved in $O(d^3)$ by the Hungarian algorithm (Kuhn, 1955; Martello & Toth, 1987). Additionally, to solve the optimization problem in Equation (8), Algorithm 1 requires the complexity of $O(d^3)$ for eigendecomposition. Hence, the complexity of our proposed full parameter matching algorithm of a transformer is $O(d^3 LH)$, similar to the complexity of the feedforward. It is worth noting that the complexity can be further reduced in two ways. The first way is to match a subset of layers instead of all. In this way, the complexity of our proposed parameter matching algorithm is $O(d^3 L_s H)$ where L_s denotes the number of selected layers. The second way is to match each unit module (a single feedforward network or a single attention layer) in parallel. The decoupling of matching different modules makes it easy to implement multiprocessing. Consequently, the overall complexity becomes to $O(d^3 LH/p)$ where p is the number of processes in parallel.

Table 2. Experimental results of ViT merging over the image classification task.

	w/o match	w/o rescale	match
Simple	7.60	10.22	10.19
Fisher	17.96	18.61	18.58
Regmean	14.24	15.31	15.35
OT-ACTS-EMD	32.08	32.50	32.53
OT-ACTS	61.15	61.23	61.25
OT-WTS	57.11	57.16	57.17

5. Experiments

In this section, we conduct experiments to evaluate the effectiveness of our proposed parameter matching algorithm based on rotation symmetry. Specifically, we aim to answer the following research questions: RQ1: Can our proposed parameter matching algorithm enhance the performance of model fusion for transformer-based models? RQ2: How does rotation symmetry contribute to parameter matching for self-attention layers? RQ3: How does parameter matching influence the loss landscape between independently trained models? RQ4: As a plugin module, does our algorithm introduce significant additional computational overhead? RQ5: Is matching all transformer layers equally important? Can we improve efficiency without compromising utility by matching only a subset of layers?

5.1. Experimental Settings

Platform. Our implementation is based on Python 3.10 and Pytorch 1.13. All fine-tuning, parameter matching, and model merging processes are conducted on a cluster equipped with Nvidia A100 80GB GPUs.

Models. To evaluate the effectiveness of our approach, we use two widely adopted transformer models: RoBERTa (Liu et al., 2019) and DeBERTa (He et al., 2021). We obtain pretrained models of RoBERTa (roberta-base with 12 attention layers) and DeBERTa-Large (microsoft/deberta-v3-large with 24 attention layers) from the Hugging Face library. For vision transformers (ViTs) (Dosovitskiy et al., 2021), we directly use the pretrained models following (Imfeld et al., 2024). By selecting these three type of models, we assess the effectiveness of our method across different model scales and downstream tasks.

Finetuning and Matching. In the experiments, each pretrained language model is fine-tuned for 20 epochs on each dataset individually. We set the learning rate at $1e-5$, the batch size at 16, and the warmup ratio at 0.06 for each model. After fine-tuning, we perform parameter matching and model merging, considering both in-domain (pairwise

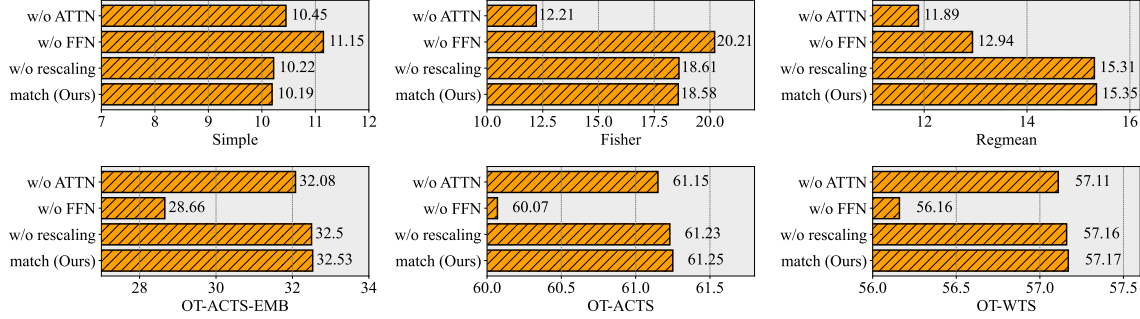


Figure 3. Ablation Study of ViT merging over the image classification task. “ATTN” is short for “attention”. We compare our matching algorithm with its three variants (w/o ATTN/FFN/rescaling) and the original performance (w/o match) on all six merging baselines.

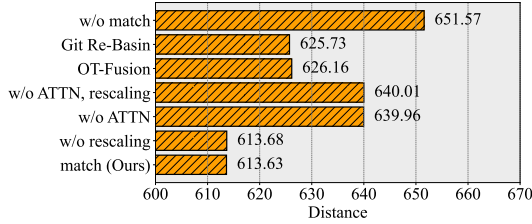


Figure 4. The Euclidean Distance of end ViT models after different parameter matching algorithms.

fine-tuned models) and out-of-domain (grouped models) experiments. Notably, we match only the parameters within the attention layers, while the classifier module is directly copied from the model fine-tuned on the corresponding downstream task. For additional details on datasets and baseline methods, please refer to Appendix C.

5.2. Performance of Model Fusion

To answer RQ1, we compare the performance of three model fusion baselines with and without our proposed parameter matching algorithm.

In-Domain Settings. We evaluate model fusion performance on Emotion Classification and Named Entity Recognition (NER) tasks, where our approach (**+match**) is integrated as a plugin module for three model fusion baselines. For each task, we fine-tune the language models on each in-domain dataset (5 for Emotion, 6 for NER) respectively and merge them pairwise. The first two columns in Table 1 present the average macro F1-score (for Emotion) and micro F1-score (for NER) of the merged models following (Jin et al., 2023).

Out-of-Domain Settings. To assess the generalization ability, we merge models trained on all in-domain NER datasets and evaluate their performance on CoNLL datasets, which serve as out-of-domain (OOD) test sets. The third column in Table 1 reports the OOD performance of the merged models.

ViT Settings. We employ three OT (optimal transportation)-based methods (OT-ACTS-EMD, OT-ACTS, and OT-WTS) from (Imfeld et al., 2024) as additional baselines. Two pretrained models are merged, and we evaluate their performance on the CIFAR-10 (Krizhevsky et al., 2009) dataset. Table 2 presents the classification accuracy of the merged models. To improve the flexibility of parameter matching, we match each layer separately before conducting model fusion. The merged model achieving the highest validation performance is selected as the final test model.

From the results in Table 1 and Table 2, we make the following observations: (1) Our parameter matching algorithm consistently improves the performance of different model fusion methods. The reason is that parameter matching helps align distant models, bringing them closer in parameter space. As a result, the merged model is more likely to approach an overall minimum. (2) Compared to RoBERTa-base, our method brings larger improvement for DeBERTa-large, suggesting that larger models benefit more from parameter matching. We explain this as larger models have more parameters that can be aligned, and in high-dimensional spaces, models tend to be more spread out. Subsequently, parameter matching helps bridge this gap more effectively. (3) Among all the fusion methods, simple fusion shows the largest improvement after parameter matching. This is likely because simple fusion directly averages model weights, so reducing the distance between model parameters brings a clear benefit. In contrast, advanced fusion methods work by aligning outputs of the end models based on input data, which is equivalent to weighted averaging. Unlike direct averaging, these methods benefit from parameter matching in a more implicit way.

5.3. Ablation Study

To answer RQ2, we conduct ablation studies to evaluate the contributions of rotation symmetry in matching different model components. These experiments are performed on the same pair of pretrained ViT models used in the previous

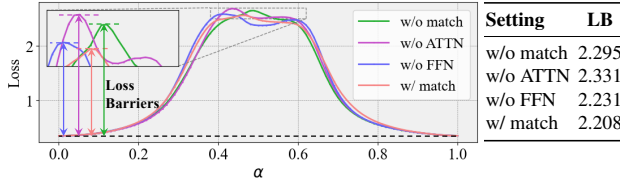


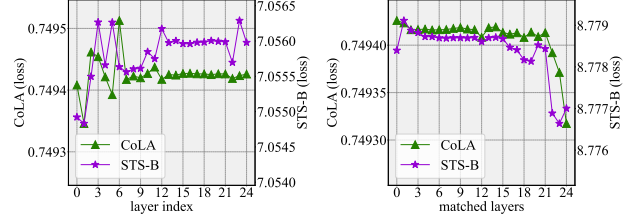
Figure 5. Loss landscapes and barriers between the two pretrained ViT models under four distinct matching settings. “LB” is short for “Loss Barrier”.

Table 3. The average runtime in seconds of the fine-tuning (top), matching (middle), and merging (bottom) stage.

	Deberta	Roberta	ViT
Fine-tuning	12983.71	3071.49	-
Matching	1.59	1.71	3.48
Simple Merging	0.13	0.09	0.22
Fisher Merging	197.47	69.57	83.67
Regmean Merging	137.67	36.44	71.02

section. We design three ablation settings: (1) *w/o ATTN*: Matches only the FFN modules, excluding self-attention modules. (2) *w/o FFN*: Matches only the self-attention modules, excluding FFN modules. (3) *w/o rescaling*: Disables the rescaling symmetry in our algorithm. Ablation results in Figure 3 demonstrate the following observations: (1) In Fisher and RegMean, attention matching contributes more significantly to the fusion process. (2) In contrast, OTFusion benefits more from FFN matching. These distinctions demonstrate the nuanced interactions between the fusion mechanisms and corresponding component alignment strategies. (3) *w/o rescaling* results in performance degradation across most settings, validating the effectiveness of rescaling symmetry integrated with rotation symmetry.

To further quantify the influence of rotation symmetry, we measure the Euclidean distance between the matched model and the anchor model in the parameter space. Instead of aggregating distances at the layer or module level, we analyze the entire model holistically to provide a more comprehensive assessment of alignment. We also compare our method against two external baselines: Git Re-Basin (Ainsworth et al., 2023) and OT-Fusion (Imfeld et al., 2024). Following the original paper, OT-Fusion is applied only to FFN modules. The results in Figure 4 indicate that our rotation symmetry-based parameter matching algorithm consistently reduces the distance between end models more effectively than permutation symmetry-based methods. This improvement can be attributed to the continuous nature of rotation symmetry, which allows for smoother and more precise alignment, especially in self-attention layers. Furthermore, the integration of rescaling symmetry leads to additional distance reductions of end models, further enhancing parameter alignment.



(a) Single-Layer Matching (b) Tail-Layers Matching

Figure 6. Evaluation loss of matching a subset of layers.

5.4. Loss Landscape Study

To address RQ3, we examine the loss landscape between the two pretrained ViT models introduced in Section 5.2, both with and without parameter matching. Specifically, we first analyze the cross entropy loss values $\mathcal{L}(\theta(\alpha))$ along the linear interpolation path between the end model parameters $\theta(\alpha) = \alpha \cdot \theta_A + (1 - \alpha) \cdot \theta_B$, $\alpha \in [0, 1]$, where θ_A and θ_B denote the end models. Next, we separately apply attention-only matching (*w/o FFN*), FFN-only matching (*w/o ATTN*), and complete matching (*w/ match*) to θ_B before constructing the same linear interpolation path with respect to θ_A . The results in Figure 5 give rise to the following observations: (1) All model pairs exhibit positive loss barriers, i.e., the range of loss value along the interpolation path, highlighting the strong inherent non-convexity of ViTs (Park & Kim, 2022). (2) Complete matching yields the lowest loss barrier, indicating that aligning parameters through rotation symmetry leads to smoother connectivity between model pairs. (3) While FFN-only matching (*w/o ATTN*) has been proven effective for MLPs and CNNs (Ainsworth et al., 2023), its effectiveness remains limited for Transformers. Their intricate loss landscape brings additional challenges for parameter matching.

5.5. Complexity Study

To answer RQ4, we evaluate the computational overhead introduced by our parameter matching algorithm. Specifically, we measure the average runtime for fine-tuning (per dataset), matching (per model pair), and merging methods (per model pair) from the main experiments, as shown in Table 3. For ViTs, we only use pre-trained models following (Imfeld et al., 2024) without fine-tuning. The results show that our parameter matching module incurs an overhead of less than 5% compared to Fisher / RegMean merging across all models, with negligible impact on overall complexity.

5.6. Matching a Subset of Layers

To answer RQ5, we investigate the effect of matching different subsets of layers. We fine-tune DeBERTa models on CoLA and STS-B from the GLUE benchmark and evaluate the performance of matched models under different subsets

of layers. Specifically, we cross-validate the importance of each attention layer in matching through two settings:

Single-Layer Matching. In this setting, we match only a single attention layer while leaving all other layers unchanged. The evaluation loss corresponding to different matched layer indices is shown in Figure 6(a).

Tail-Layers Matching. Here, we match a certain number of trailing attention layers (*e.g.*, when matching three attention layers, we align only the last three layers). The evaluation loss for different numbers of matched layers is presented in Figure 6(b).

The experimental results in Figure 6 demonstrate that matching head layers yields greater improvements in model utility compared to tail layers. Notably, the loss value drops sharply when the first 5 layers are matched. Based on this observation, we can further improve efficiency by dropping tail layers without significantly compromising utility.

6. Related Work

Parameter Space Symmetry. Parameter space symmetry refers to a set of models with different parameter values but functionally equivalent. This concept has been extensively studied in the context of deep neural networks, as it plays a crucial role in understanding model behavior and training dynamics. Examples of parameter space symmetries include rescaling symmetry (Neyshabur et al., 2015; Badrinarayanan et al., 2015; Du et al., 2018; Meng et al., 2019), scaling symmetry (Kunin et al., 2021), and translation symmetry (Kunin et al., 2021). These symmetries have been identified in conventional deep neural networks to provide deeper insights into training dynamics and to accelerate the optimization process (Zhao et al., 2022; 2023; 2024). Another important type of parameter space symmetry is permutation symmetry, which has been shown to closely relate to the manifold of global minima and critical points (Fukumizu & Amari, 2000; Brea et al., 2019; Simsek et al., 2021; Benton et al., 2021; Entezari et al., 2022; Ainsworth et al., 2023). The permutation symmetry can also be used to align (match) the outputs or model parameters of different end models with the same architecture (Singh & Jaggi, 2020; Wang et al., 2020; Ainsworth et al., 2023; Peña et al., 2023; Imfeld et al., 2024; Navon et al., 2024). Complementary to parameter space symmetry, recent advances in neural functionals and metanetworks (Navon et al., 2023; Zhou et al., 2023b;a; Lim et al., 2024a; Tran et al., 2024) explore permutation-equivariant functionals that operate directly on model weights across diverse architectures. Some concurrent works (Liu, 2024; Tran et al., 2024) investigate similar forms of rotation symmetry in neural networks. Liu (2024) shows that the mirror symmetry leads

to low-rankness. Meanwhile, Tran et al. (2024) leverages the rotation symmetry to construct transformer-based neural functional networks. In comparison, our study focuses on the role of rotation symmetry in model fusion and proposes a theoretically optimal parameter matching approach based on the properties of rotation symmetries.

Model Fusion. The goal of model fusion (Li et al., 2023) is to merge multiple available end models (with the same architecture) to obtain a stronger model. The scenarios of model fusion can be flexible. When training on the same dataset, model fusion can be used to improve the model utility or generalization by merging models trained with different configurations or in different stages (Izmailov et al., 2018; Gupta et al., 2020; Cha et al., 2021; Wortsman et al., 2022; Rame et al., 2022; Arpit et al., 2022; Huang et al., 2024; Yadav et al., 2024; Hammoud et al., 2024). As a representative method in this setting, ModelSoup (Wortsman et al., 2022) greedily averages the models fine-tuned with different hyperparameter configurations to improve the utility and robustness of the model. In addition, when training on different datasets or tasks, model fusion can be used to improve out-of-domain generalization or multitasking of the model (Matena & Raffel, 2022; Choshen et al., 2022; Li et al., 2022; Jin et al., 2023; Zhou et al., 2024), especially for language models. A state-of-the-art merging algorithm, RegMean (Jin et al., 2023), successfully merges language models fine-tuned over different tasks and improves the model’s out-of-distribution generalization. Moreover, model fusion plays a pivotal role in federated learning (Konečný et al., 2016; McMahan et al., 2017; Wang et al., 2020) when the local updates are collected to make a global update. FedAvg (McMahan et al., 2017) is a classical merging algorithm that directly computes the average of the local models as the updated global model. Recent studies propose to incorporate the permutation symmetry to align the neurons of different end models (Wang et al., 2020; Singh & Jaggi, 2020; Ainsworth et al., 2023). However, these methods fail to achieve a desirable performance when tackling transformer-based models (Jin et al., 2023).

7. Conclusion

In this paper, we introduced rotation symmetry as a novel type of parameter space symmetry for transformers, extending the concept of permutation symmetry to continuous spaces. Building on this foundation, we proposed a theoretically optimal parameter matching algorithm to enhance the fusion of transformer models in a plug-and-play manner. To validate our approach, we conducted extensive experiments on real-world NLP and vision benchmarks. The results demonstrated that incorporating rotation symmetry effectively and efficiently facilitates transformer model fusion, showcasing our method’s practical utility.

Impact Statement

Our study provides novel insights into the parameter space symmetry of transformers and establishes a practical framework for advancing model fusion techniques. Bridging theoretical innovations with practical applications, this work reveals the potential for leveraging parameter space symmetries in deep learning research. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

Acknowledgment

This work is supported in part by the National Science Foundation (NSF) under grants IIS-2006844, IIS-2144209, IIS-2223769, CNS-2154962, BCS-2228534, and CMMI-2411248; the Office of Naval Research (ONR) under grant N000142412636; and the Commonwealth Cyber Initiative (CCI) under grant VV-1Q24-011.

References

- Ainsworth, S., Hayase, J., and Srinivasa, S. Git re-basin: Merging models modulo permutation symmetries. In *International Conference on Learning Representations*, 2023.
- Armenta, M. and Jodoin, P.-M. The representation theory of neural networks. *Mathematics*, 9(24):3216, 2021.
- Arpit, D., Wang, H., Zhou, Y., and Xiong, C. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *Advances in Neural Information Processing Systems*, 35:8265–8277, 2022.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Badrinarayanan, V., Mishra, B., and Cipolla, R. Symmetry-invariant optimization in deep networks. *arXiv preprint arXiv:1511.01754*, 2015.
- Benton, G., Maddox, W., Lotfi, S., and Wilson, A. G. G. Loss surface simplexes for mode connecting volumes and fast ensembling. In *International Conference on Machine Learning*, pp. 769–779, 2021.
- Bostan, L.-A.-M. and Klinger, R. An analysis of annotated corpora for emotion classification in text. In Bender, E. M., Derczynski, L., and Isabelle, P. (eds.), *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 2104–2119, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://aclanthology.org/C18-1179>.
- Brea, J., Simsek, B., Illing, B., and Gerstner, W. Weight-space symmetry in deep networks gives rise to permutation saddles, connected by equal-loss valleys across the loss landscape. *arXiv preprint arXiv:1907.02911*, 2019.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.
- Burkard, R. E. and Cela, E. Linear assignment problems and extensions. In *Handbook of combinatorial optimization: Supplement volume A*, pp. 75–149. 1999.
- Cha, J., Chun, S., Lee, K., Cho, H.-C., Park, S., Lee, Y., and Park, S. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021.
- Choshen, L., Venezian, E., Slonim, N., and Katz, Y. Fusing finetuned models for better pretraining. *arXiv preprint arXiv:2204.03044*, 2022.
- Clark, K., Luong, M.-T., Le, Q. V., and Manning, C. D. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*, 2020.
- Daheim, N., Möllenhoff, T., Ponti, E., Gurevych, I., and Khan, M. E. Model merging by uncertainty-based gradient matching. In *International Conference on Learning Representations*, 2024.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- Dietterich, T. G. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pp. 1–15, 2000.
- Dong, X., Yu, Z., Cao, W., Shi, Y., and Ma, Q. A survey on ensemble learning. *Frontiers of Computer Science*, 14: 241–258, 2020.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houselby, N. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. URL <https://arxiv.org/abs/2010.11929>.
- Du, S. S., Hu, W., and Lee, J. D. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *Advances in neural information processing systems*, 31, 2018.

- Entezari, R., Sedghi, H., Saukh, O., and Neyshabur, B. The role of permutation invariance in linear mode connectivity of neural networks. In *International Conference on Learning Representations*, 2022.
- Ferbach, D., Goujaud, B., Gidel, G., and Dieuleveut, A. Proving linear mode connectivity of neural networks via optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pp. 3853–3861, 2024.
- Fukumizu, K. and Amari, S.-i. Local minima and plateaus in hierarchical structures of multilayer perceptrons. *Neural networks*, 13(3):317–327, 2000.
- Glorot, X., Bordes, A., and Bengio, Y. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 315–323, 2011.
- Godfrey, C., Brown, D., Emerson, T., and Kvinge, H. On the symmetries of deep learning models and their internal representations. *Advances in Neural Information Processing Systems*, 35:11893–11905, 2022.
- Gower, J. C. and Dijksterhuis, G. B. *Procrustes problems*, volume 30. OUP Oxford, 2004.
- Grigsby, E., Lindsey, K., and Rolnick, D. Hidden symmetries of relu networks. In *International Conference on Machine Learning*, pp. 11734–11760, 2023.
- Gupta, V., Serrano, S. A., and DeCoste, D. Stochastic weight averaging in parallel: Large-batch training that generalizes well. In *International Conference on Learning Representations*, 2020.
- Hammoud, H., Michieli, U., Pizzati, F., Torr, P., Bibi, A., Ghanem, B., and Ozay, M. Model merging and safety alignment: One bad model spoils the bunch. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 13033–13046, 2024.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- He, P., Liu, X., Gao, J., and Chen, W. DeBERTa: Decoding-enhanced bert with disentangled attention, 2021.
- He, Y., Zheng, Z., Soga, P., Zhu, Y., Dong, Y., and Li, J. Explaining graph neural networks with large language models: A counterfactual perspective on molecule graphs. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 7079–7096, 2024.
- Hecht-Nielsen, R. On the algebraic structure of feedforward network weight spaces. In *Advanced Neural Computers*, pp. 129–135. 1990.
- Hovy, E., Marcus, M., Palmer, M., Ramshaw, L., and Weischedel, R. OntoNotes: The 90% solution. In Moore, R. C., Bilmes, J., Chu-Carroll, J., and Sanderson, M. (eds.), *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pp. 57–60, New York City, USA, June 2006. Association for Computational Linguistics. URL <https://aclanthology.org/N06-2015>.
- Huang, C., Ye, P., Chen, T., He, T., Yue, X., and Ouyang, W. Emr-merging: Tuning-free high-performance model merging. *Advances in Neural Information Processing Systems*, 37:122741–122769, 2024.
- Imfeld, M., Graldi, J., Giordano, M., Hofmann, T., Anagnostidis, S., and Singh, S. P. Transformer fusion with optimal transport. In *International Conference on Learning Representations*, 2024.
- Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., and Wilson, A. G. Averaging weights leads to wider optima and better generalization. In *34th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 876–885, 2018.
- Jin, X., Ren, X., Preotiuc-Pietro, D., and Cheng, P. Data-less knowledge fusion by merging weights of language models. In *International Conference on Learning Representations*, 2023.
- Kabsch, W. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976.
- Kalogeropoulos, I., Bouritsas, G., and Panagakis, Y. Scale equivariant graph metanetworks. *Advances in neural information processing systems*, 37:106800–106840, 2024.
- Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., and Bacon, D. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Kuhn, H. W. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- Kunin, D., Sagastuy-Brena, J., Ganguli, S., Yamins, D. L., and Tanaka, H. Neural mechanics: Symmetry and broken conservation laws in deep learning dynamics. In *International Conference on Learning Representations*, 2021.

- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880, 2020.
- Li, M., Gururangan, S., Dettmers, T., Lewis, M., Althoff, T., Smith, N. A., and Zettlemoyer, L. Branch-train-merge: Embarrassingly parallel training of expert language models. In *First Workshop on Interpolation Regularizers and Beyond at NeurIPS 2022*, 2022.
- Li, W., Peng, Y., Zhang, M., Ding, L., Hu, H., and Shen, L. Deep model fusion: A survey. *arXiv preprint arXiv:2309.15698*, 2023.
- Lim, D., Maron, H., Law, M. T., Lorraine, J., and Lucas, J. Graph metanetworks for processing diverse neural architectures. In *International Conference on Learning Representations*, 2024a.
- Lim, D., Putterman, T., Walters, R., Maron, H., and Jegelka, S. The empirical impact of neural parameter symmetries, or lack thereof. *Advances in Neural Information Processing Systems*, 37:28322–28358, 2024b.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Liu, Z. Symmetry induces structure and constraint of learning. In *International Conference on Machine Learning*, 2024.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.
- Lubana, E. S., Bigelow, E. J., Dick, R. P., Krueger, D., and Tanaka, H. Mechanistic mode connectivity. In *International Conference on Machine Learning*, pp. 22965–23004, 2023.
- Martello, S. and Toth, P. Linear assignment problems. In *North-Holland Mathematics Studies*, volume 132, pp. 259–282. 1987.
- Matena, M. S. and Raffel, C. A. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282, 2017.
- Meng, Q., Zheng, S., Zhang, H., Chen, W., Ye, Q., Ma, Z.-M., Yu, N., and Liu, T.-Y. G-sgd: Optimizing relu neural networks in its positively scale-invariant space. In *International Conference on Learning Representations*, 2019.
- Nair, V. and Hinton, G. E. Rectified linear units improve restricted boltzmann machines. In *International Conference on Machine Learning*, pp. 807–814, 2010.
- Navon, A., Shamsian, A., Achituve, I., Fetaya, E., Chechik, G., and Maron, H. Equivariant architectures for learning in deep weight spaces. In *International Conference on Machine Learning*, pp. 25790–25816, 2023.
- Navon, A., Shamsian, A., Fetaya, E., Chechik, G., Dym, N., and Maron, H. Equivariant deep weight space alignment. In *International Conference on Machine Learning*, pp. 37376–37395, 2024.
- Neyshabur, B., Salakhutdinov, R. R., and Srebro, N. Path-sgd: Path-normalized optimization in deep neural networks. *Advances in neural information processing systems*, 28, 2015.
- Park, N. and Kim, S. How do vision transformers work?, 2022. URL <https://arxiv.org/abs/2202.06709>.
- Peña, F. A. G., Medeiros, H. R., Dubail, T., Aminbeidokhti, M., Granger, E., and Pedersoli, M. Re-basin via implicit sinkhorn differentiation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20237–20246, 2023.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763, 2021.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21 (140):1–67, 2020.
- Rame, A., Kirchmeyer, M., Rahier, T., Rakotomamonjy, A., Gallinari, P., and Cord, M. Diverse weight averaging for

- out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 35:10821–10836, 2022.
- Rossi, S., Singh, A., and Hannagan, T. On permutation symmetries in bayesian neural network posteriors: a variational perspective. *Advances in Neural Information Processing Systems*, 36, 2023.
- Sagi, O. and Rokach, L. Ensemble learning: A survey. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 8(4):e1249, 2018.
- Schönemann, P. H. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966.
- Simsek, B., Ged, F., Jacot, A., Spadaro, F., Hongler, C., Gerstner, W., and Brea, J. Geometry of the loss landscape in overparameterized neural networks: Symmetries and invariances. In *International Conference on Machine Learning*, pp. 9722–9732, 2021.
- Singh, S. P. and Jaggi, M. Model fusion via optimal transport. *Advances in Neural Information Processing Systems*, 33:22045–22055, 2020.
- Strang, G. *Linear algebra and its applications*. Academic Press, 1976.
- Tatro, N., Chen, P.-Y., Das, P., Melnyk, I., Sattigeri, P., and Lai, R. Optimizing mode connectivity via neuron alignment. *Advances in Neural Information Processing Systems*, 33:15300–15311, 2020.
- Tjong Kim Sang, E. F. and De Meulder, F. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pp. 142–147, 2003. URL <https://aclanthology.org/W03-0419>.
- Tran, V.-H., Vo, T. N., The, A. N., Huu, T. T., Nguyen-Nhat, M.-K., Tran, T., Pham, D.-T., and Nguyen, T. M. Equivariant neural functional networks for transformers. *arXiv preprint arXiv:2410.04209*, 2024.
- Umeyama, S. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13(04):376–380, 1991.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. Glue: A multi-task benchmark and analysis platform for natural language understanding, 2019.
- Wang, H., Yurochkin, M., Sun, Y., Papailiopoulos, D., and Khazaeni, Y. Federated learning with matched averaging. In *International Conference on Learning Representations*, 2020.
- Wei, X. and Liu, L. Are large language models good in-context learners for financial sentiment analysis? *arXiv preprint arXiv:2503.04873*, 2025.
- 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, 2022.
- Yadav, P., Tam, D., Choshen, L., Raffel, C. A., and Bansal, M. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36, 2023.
- Yadav, P., Vu, T., Lai, J., Chronopoulou, A., Faruqui, M., Bansal, M., and Munkhdalai, T. What matters for model merging at scale? *arXiv preprint arXiv:2410.03617*, 2024.
- Yang, E., Wang, Z., Shen, L., Liu, S., Guo, G., Wang, X., and Tao, D. Adamerging: Adaptive model merging for multi-task learning. In *International Conference on Learning Representations*, 2024.
- Yun, S., Jeong, M., Kim, R., Kang, J., and Kim, H. J. Graph transformer networks. *Advances in neural information processing systems*, 32, 2019.
- Yurochkin, M., Agarwal, M., Ghosh, S., Greenewald, K., Hoang, N., and Khazaeni, Y. Bayesian nonparametric federated learning of neural networks. In *International conference on machine learning*, pp. 7252–7261, 2019.
- Zamir, G., Dokania, A., Zhao, B., and Yu, R. Improving learning to optimize using parameter symmetries. *arXiv preprint arXiv:2504.15399*, 2025.
- Zhang, B., Chen, Z., Zheng, Z., Li, J., and Chen, H. Resolving editing-unlearning conflicts: A knowledge codebook framework for large language model updating. *arXiv preprint arXiv:2502.00158*, 2025.
- Zhao, B., Dehmamy, N., Walters, R., and Yu, R. Symmetry teleportation for accelerated optimization. *Advances in neural information processing systems*, 35:16679–16690, 2022.
- Zhao, B., Ganev, I., Walters, R., Yu, R., and Dehmamy, N. Symmetries, flat minima, and the conserved quantities of gradient flow. In *The Eleventh International Conference on Learning Representations*, 2023.

- Zhao, B., Gower, R. M., Walters, R., and Yu, R. Improving convergence and generalization using parameter symmetries. In *International Conference on Learning Representations*, 2024.
- Zheng, Z., Dong, Y., Wang, S., Liu, H., Wang, Q., and Li, J. Kg-cf: Knowledge graph completion with context filtering under the guidance of large language models. In *2024 IEEE International Conference on Big Data (BigData)*, pp. 805–810, 2024.
- Zhou, A., Yang, K., Burns, K., Cardace, A., Jiang, Y., Sokota, S., Kolter, J. Z., and Finn, C. Permutation equivariant neural functionals. *Advances in neural information processing systems*, 36:24966–24992, 2023a.
- Zhou, A., Yang, K., Jiang, Y., Burns, K., Xu, W., Sokota, S., Kolter, J. Z., and Finn, C. Neural functional transformers. *Advances in neural information processing systems*, 36: 77485–77502, 2023b.
- Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., and Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 11106–11115, 2021.
- Zhou, Y., Song, L., Wang, B., and Chen, W. Metagpt: Merging large language models using model exclusive task arithmetic. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 1711–1724, 2024.
- Zhou, Z., Yang, Y., Yang, X., Yan, J., and Hu, W. Going beyond linear mode connectivity: The layerwise linear feature connectivity. *Advances in Neural Information Processing Systems*, 36, 2023c.
- Zhu, Y., Wu, L., Zhang, B., Wang, S., Guo, Q., Hong, L., Simon, L., and Li, J. Understanding and modeling job marketplace with pretrained language models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 5143–5150, 2024.
- Ziyin, L., Wang, M., Li, H., and Wu, L. Parameter symmetry and noise equilibrium of stochastic gradient descent. *arXiv preprint arXiv:2402.07193*, 2024.
- Ziyin, L., Xu, Y., and Chuang, I. L. Remove symmetries to control model expressivity and improve optimization. In *International Conference on Learning Representations*, 2025.

A. Proof

Theorem 4.1. *The following optimization problem has a closed-form solution.*

$$\min_{\mathbf{R}_1, \mathbf{R}_2 \in \mathcal{R}} \left\| \begin{bmatrix} \mathbf{W}_{Q_1}^\top & \mathbf{W}_{Q_2}^\top \\ \mathbf{b}_{Q_1} & \mathbf{b}_{Q_2} \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 \\ -\mathbf{R}_2 \end{bmatrix} \right\|_F^2 + \left\| \begin{bmatrix} \mathbf{W}_{K_1}^\top & \mathbf{W}_{K_2}^\top \\ \mathbf{b}_{K_1} & \mathbf{b}_{K_2} \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 \\ -\mathbf{R}_2 \end{bmatrix} \right\|_F^2. \quad (13)$$

The solution is given by

$$\mathbf{R}_1 = \mathbf{U}\mathbf{V}^\top, \mathbf{R}_2 = \mathbf{I}, \quad (14)$$

where \mathbf{I} is the identity matrix and $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \mathbf{W}_{Q_1}\mathbf{W}_{Q_2}^\top + \mathbf{W}_{K_1}\mathbf{W}_{K_2}^\top + \mathbf{b}_{Q_1}^\top\mathbf{b}_{Q_2} + \mathbf{b}_{K_1}^\top\mathbf{b}_{K_2}$ is the result of eigendecomposition.

Proof. We first show the process of converting the optimization problem to an Orthogonal Procrustes problem (Schönmann, 1966; Gower & Dijksterhuis, 2004). We can obtain that the optimization problem is equivalent to the following one.

$$\min_{\mathbf{R}_1, \mathbf{R}_2 \in \mathcal{R}} \left\| \mathbf{W}_{Q_1}^\top \mathbf{R}_1 - \mathbf{W}_{Q_2}^\top \mathbf{R}_2 \right\|_F^2 + \left\| \mathbf{b}_{Q_1} \mathbf{R}_1 - \mathbf{b}_{Q_2} \mathbf{R}_2 \right\|_F^2 + \left\| \mathbf{W}_{K_1}^\top \mathbf{R}_1 - \mathbf{W}_{K_2}^\top \mathbf{R}_2 \right\|_F^2 + \left\| \mathbf{b}_{K_1} \mathbf{R}_1 - \mathbf{b}_{K_2} \mathbf{R}_2 \right\|_F^2. \quad (15)$$

Considering that $\mathbf{R}_1, \mathbf{R}_2 \in \mathcal{R}$ and $\mathbf{R}_1, \mathbf{R}_2$ are nonsingular matrices, we let $\mathbf{R} = \mathbf{R}_1 \mathbf{R}_2^{-1}$ and convert Equation (15) to the following one:

$$\min_{\mathbf{R}, \mathbf{R}_2 \in \mathcal{R}} \left\| (\mathbf{W}_{Q_1}^\top \mathbf{R} - \mathbf{W}_{Q_2}^\top) \mathbf{R}_2 \right\|_F^2 + \left\| (\mathbf{b}_{Q_1} \mathbf{R} - \mathbf{b}_{Q_2}) \mathbf{R}_2 \right\|_F^2 + \left\| (\mathbf{W}_{K_1}^\top \mathbf{R} - \mathbf{W}_{K_2}^\top) \mathbf{R}_2 \right\|_F^2 + \left\| (\mathbf{b}_{K_1} \mathbf{R} - \mathbf{b}_{K_2}) \mathbf{R}_2 \right\|_F^2. \quad (16)$$

As multiplying \mathbf{R}_2 preserves the Frobenius norm, we can remove the \mathbf{R}_2 terms from the objective and obtain an Orthogonal Procrustes problem as follows:

$$\min_{\mathbf{R} \in \mathcal{R}} \left\| \mathbf{W}_{Q_1}^\top \mathbf{R} - \mathbf{W}_{Q_2}^\top \right\|_F^2 + \left\| \mathbf{b}_{Q_1} \mathbf{R} - \mathbf{b}_{Q_2} \right\|_F^2 + \left\| \mathbf{W}_{K_1}^\top \mathbf{R} - \mathbf{W}_{K_2}^\top \right\|_F^2 + \left\| \mathbf{b}_{K_1} \mathbf{R} - \mathbf{b}_{K_2} \right\|_F^2. \quad (17)$$

We take a look at the first term $\left\| \mathbf{W}_{Q_1}^\top \mathbf{R} - \mathbf{W}_{Q_2}^\top \right\|_F^2$ and have:

$$\begin{aligned} & \min_{\mathbf{R} \in \mathcal{R}} \left\| \mathbf{W}_{Q_1}^\top \mathbf{R} - \mathbf{W}_{Q_2}^\top \right\|_F^2 \\ &= \min_{\mathbf{R} \in \mathcal{R}} \langle \mathbf{W}_{Q_1}^\top \mathbf{R} - \mathbf{W}_{Q_2}^\top, \mathbf{W}_{Q_1}^\top \mathbf{R} - \mathbf{W}_{Q_2}^\top \rangle_F \\ &= \max_{\mathbf{R} \in \mathcal{R}} \langle \mathbf{W}_{Q_1}^\top \mathbf{R}, \mathbf{W}_{Q_2}^\top \rangle_F \\ &= \max_{\mathbf{R} \in \mathcal{R}} \text{tr}(\mathbf{R}^\top \mathbf{W}_{Q_1} \mathbf{W}_{Q_2}^\top) \\ &= \max_{\mathbf{R} \in \mathcal{R}} \langle \mathbf{R}, \mathbf{W}_{Q_1} \mathbf{W}_{Q_2}^\top \rangle_F. \end{aligned} \quad (18)$$

Similarly, we can convert Equation (17) to the following one:

$$\max_{\mathbf{R} \in \mathcal{R}} \langle \mathbf{R}, \mathbf{W}_{Q_1} \mathbf{W}_{Q_2}^\top + \mathbf{W}_{K_1} \mathbf{W}_{K_2}^\top + \mathbf{b}_{Q_1}^\top \mathbf{b}_{Q_2} + \mathbf{b}_{K_1}^\top \mathbf{b}_{K_2} \rangle_F. \quad (19)$$

We then conduct the singular value decomposition to the matrix $\mathbf{W}_{Q_1} \mathbf{W}_{Q_2}^\top + \mathbf{W}_{K_1} \mathbf{W}_{K_2}^\top + \mathbf{b}_{Q_1}^\top \mathbf{b}_{Q_2} + \mathbf{b}_{K_1}^\top \mathbf{b}_{K_2} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ and have:

$$\begin{aligned} & \max_{\mathbf{R} \in \mathcal{R}} \langle \mathbf{R}, \mathbf{W}_{Q_1} \mathbf{W}_{Q_2}^\top + \mathbf{W}_{K_1} \mathbf{W}_{K_2}^\top + \mathbf{b}_{Q_1}^\top \mathbf{b}_{Q_2} + \mathbf{b}_{K_1}^\top \mathbf{b}_{K_2} \rangle_F \\ &= \max_{\mathbf{R} \in \mathcal{R}} \text{tr}(\mathbf{V}\mathbf{\Sigma}\mathbf{U}^\top \mathbf{R}) \\ &= \max_{\mathbf{R} \in \mathcal{R}} \text{tr}(\mathbf{\Sigma}\mathbf{U}^\top \mathbf{R}\mathbf{V}) \\ &= \text{tr}(\mathbf{\Sigma}), \end{aligned} \quad (20)$$

where $\mathbf{U}^\top \mathbf{R}\mathbf{V} = \mathbf{I}$, i.e., $\mathbf{R} = \mathbf{U}\mathbf{V}^\top$. Finally, we incorporate $\mathbf{R} = \mathbf{U}\mathbf{V}^\top$ into $\mathbf{R} = \mathbf{R}_1 \mathbf{R}_2^{-1}$ and let $\mathbf{R}_2 = \mathbf{I}$ for simplicity. Subsequently, we have $\mathbf{R}_1 = \mathbf{U}\mathbf{V}^\top$ and $\mathbf{R}_2 = \mathbf{I}$ as a closed-form solution of the optimization problem in Theorem 4.1. \square

From the proof, we can observe that the optimization problem in Theorem 4.1 has infinite pairs of solutions $\{\mathbf{R}_1, \mathbf{R}_2\}$ regarding the value of \mathbf{R}_2 . In this paper, we let $\mathbf{R}_2 = \mathbf{I}$ for simplicity, which makes model 2 an anchor model. However, the value of \mathbf{R}_2 can be further adjusted to benefit some data-dependent model fusion techniques (Singh & Jaggi, 2020; Jin et al., 2023). We leave this part to future work.

B. Comparison with Related Work

Git Re-Basin (Ainsworth et al., 2023) introduces three model fusion algorithms: weight matching, activation matching, and straight-through matching, based on permutation symmetry. While these methods are effective for CNNs and MLPs, they do not easily extend to transformers due to the discrete nature of permutation symmetry. In a discrete space, the parameter matching problem is equivalent to solving a sum of bilinear assignment problems, which is NP-hard (Ainsworth et al., 2023). To address this limitation, we introduce rotation symmetry, which extends the symmetry space to a continuous domain, allowing for a closed-form solution to parameter matching. Unlike permutation symmetry, rotation symmetry provides a more flexible and efficient solution for aligning parameters in transformers.

Additionally, our study emphasizes the advantage of weight matching as a plug-and-play module for model fusion, as it preserves model functionality while reducing the inner distance of end models after matching. OT-ACTS (Imfeld et al., 2024) also includes a parameter matching step, leveraging optimal transport for model fusion. However, its matching module is restricted to permutation operations, and its fusion method is limited to simple merging. In contrast, our approach generalizes parameter matching to rotation operations and integrates it with more advanced merging strategies, demonstrating the potential of parameter matching to enhance model fusion.

A concurrent work by Tran et al. (2024) explores a similar form of rotation symmetry. However, their focus is on constructing functionally equivalent networks for transformers, whereas we focus on leveraging rotation symmetry for model fusion, proposing a theoretically optimal parameter matching algorithm specifically designed to improve fusion performance.

C. Experimental Details

Datasets. Similar with (Jin et al., 2023), We employ emotion classification and named entity recognition (NER) as the tasks for the main experiments. The emotion classification datasets are extracted from (Bostan & Klinger, 2018). Five of them are selected as in-domain datasets (emoint, ssec, electoraltweets, grounded_emotions, affectivetext), and five others are good datasets (dailydialog, crowdflower, tec, tales-emotion, isear). For the NER task, we choose OntoNotes (Hovy et al., 2006) for model finetuning and CoNLL (Tjong Kim Sang & De Meulder, 2003) for out-of-domain evaluation. Additionally, in the ablation study and subset layers matching, we employ datasets (STS-B, SST-2, CoLA) from GLUE (Wang et al., 2019) to analyze the module performance.

Baselines. We compared the performance of our method with two matching baselines:

- **OTFusion** (Singh & Jaggi, 2020) mitigates the positional mismatch of functionally analogous neurons across end models by employing optimal transport to align neurons layer-wise, thus enabling more effective model merging.
- **Git Re-Basin** (Ainsworth et al., 2023) matches the parameters of end models through learned permutation symmetries in the parameter space, aiming to transform their parameters into a shared basin within the loss landscape and facilitate seamless merging.

Meanwhile, to address RQ1, we further select six merging methods as baselines:

- **Simple** (Wortsman et al., 2022) directly averages the weight matrices of independently pretrained end models.
- **Fisher** (Matena & Raffel, 2022) approximates each end model’s posterior as a Gaussian distribution with a precision matrix derived from its Fisher information, thereby casting model merging as a Maximum Likelihood Estimation (MLE) problem over these posterior distributions.
- **RegMean** (Jin et al., 2023) formulates model merging as a regression problem, aiming to minimize the Euclidean distance between the predictions of the merged model and individual end models over a set of input samples.
- **OT-WTS** (Imfeld et al., 2024) extends OTFusion by introducing soft neuron alignment, which replaces the rigid one-to-one neuron matching with a more flexible, weighted correspondence between neurons. Subsequent model merging is then guided by these soft alignment weights.
- **OT-ACTS** and **OT-ACTS-EMD** (Imfeld et al., 2024) are two variants of **OT-WTS** that match neurons based on their

activations rather than directly aligning parameter matrices. Specifically, OT-ACTS measures activation similarity using Euclidean distance, while OT-ACTS-EMD employs Earth Mover’s Distance for a more fine-grained comparison.

The code implementation is based on (Jin et al., 2023)² for NLP tasks and (Imfeld et al., 2024)³ for vision tasks.

D. Limitations

In this study, we proposed a theoretically optimal parameter matching algorithm based on rotation symmetry. However, the optimality guarantees of our method are subject to several important prerequisites.

First, our theoretical analysis in Theorem 4.1 focuses on the binary model matching case, where the optimal solution is independent of the choice of anchor model. This property, however, does not naturally extend to the multi-model setting. When matching multiple models, a straightforward extension involves selecting a single model as an anchor and aligning the others to it using pairwise rotation-based matching. Yet, in this case, the global optimality depends on the definition of the overall distance metric and may not be preserved. An alternative is iterative pairwise matching, but this introduces path dependency, i.e., the final result can vary depending on the order of merging. Extending our framework to the optimal multi-model matching problem remains an open challenge for future work.

Second, the scope of our parameter matching algorithm is within the rotation symmetry, as defined in Equation (5). While this serves as a tractable and well-motivated instance of parameter space symmetry, we recognize that Transformers exhibit a richer set of structural symmetries beyond permutation, rotation, and rescaling. For example, recent work suggests that geometric properties of normalization and softmax layers may enlarge the symmetry set of previous model layers (Kunin et al., 2021; Lim et al., 2024a). Characterizing the complete symmetry group of transformers remains an open question.

We also considered the possibility of extending rotation matrices to invertible matrices in terms of parameter space symmetries. Although invertible matrices expand the symmetry set, they are less suitable for practical parameter matching. Specifically, our closed-form solution in Equation (9) relies on the orthogonality of the rotation matrices. If replacing them with general invertible matrices, it becomes difficult to derive a closed-form solution for Equation (9), leading to suboptimal solution and increased complexity. Additionally, parameter matching requires computation of the inverse matrix of every \mathbf{R} , which is impractical for invertible matrices. In contrast, orthogonal matrices benefit from computationally efficient inverses (via transposition) and theoretical simplicity.

Finally, we apply the rescaling symmetry after optimizing for rotation-based matching. However, our sequential approach (first optimizing rotation \mathbf{R} , then rescaling weight α) does not guarantee global optimality for the joint optimization over (\mathbf{R}, α) . Deriving a globally optimal solution would require solving Equation (12) in closed form, which remains non-trivial. Hence, we consider our method as a practical approximation to the joint optimization solution which balances theoretical soundness with computational efficiency. Exploring a fully joint optimization framework for rotation and rescaling symmetries can be a promising direction for future research.

²<https://github.com/bloomberg/dataless-model-merging>

³<https://github.com/graldij/transformer-fusion>