

How Not to Stitch Representations to Measure Similarity: Task Loss Matching Versus Direct Matching

András Balogh¹, Márk Jelasity^{1,2}

¹ University of Szeged, Hungary

² HUN-REN-SZTE Research Group on AI, Szeged, Hungary
{abalogh, jelasity}@inf.u-szeged.hu

Abstract

Measuring the similarity of the internal representations of deep neural networks is an important and challenging problem. Model stitching has been proposed as a possible approach, where two half-networks are connected by mapping the output of the first half-network to the input of the second one. The representations are considered functionally similar if the resulting stitched network achieves good task-specific performance. The mapping is normally created by training an affine stitching layer on the task at hand while freezing the two half-networks, a method called task loss matching. Here, we argue that task loss matching may be very misleading as a similarity index. For example, it can indicate very high similarity between very distant layers, whose representations are known to have different functional properties. Moreover, it can indicate very distant layers to be more similar than architecturally corresponding layers. Even more surprisingly, when comparing layers within the same network, task loss matching often indicates that some layers are more similar to a layer than itself. We argue that the main reason behind these problems is that task loss matching tends to create out-of-distribution representations to improve task-specific performance. We demonstrate that direct matching (when the mapping minimizes the distance between the stitched representations) does not suffer from these problems. We compare task loss matching, direct matching, and well-known similarity indices such as CCA and CKA. We conclude that direct matching strikes a good balance between the structural and functional requirements for a good similarity index.

Code — <https://github.com/szegedai/stitching-ood>

1 Introduction

Understanding the internal representations that emerge as a result of training machine learning models is a crucial problem, and an important aspect of this problem is to measure the similarity of representations (Kornblith et al. 2019).

However, this is a poorly defined problem, because similarity can be approached from many angles. One choice is *structural similarity* based on the geometric properties of the distribution of the activations of, for example, a set of neurons of interest (Kornblith et al. 2019; Raghu et al. 2017; Morcos, Raghu, and Bengio 2018). Another approach

is *functional similarity*, that is, to ask the question whether one representation can be transformed into the other while preserving its functionally important aspects (Csiszárík et al. 2021; Bansal, Nakkiran, and Barak 2021a).

Functional similarity has the great advantage that it grounds the meaning of similarity. A purely structural measure conveys very little information on whether the two compared representations are actually compatible or interchangeable in any practical sense (Ding, Denain, and Steinhardt 2021). However, we argue that purely functional similarity approaches are not a panacea.

1.1 Contributions

Drawbacks of functional similarity. We show that a purely functional notion of similarity also has significant drawbacks. Not taking the structural properties into account results in anomalies such as representations of one model being transformed into highly out-of-distribution representations in the other model that do perform well functionally but strongly violate some basic sanity checks required from any similarity index, such as identifying architecturally corresponding layers.

A hybrid similarity index. We argue that for measuring similarity, the most promising approach is to combine structural and functional aspects. We examine one such approach: model stitching based on the direct matching of representations. We find that it is a good compromise between purely structural and purely functional approaches.

1.2 Related Work

Structural Similarity. Several methods have been proposed to quantitatively compare the learned internal representations of neural networks based on geometrical and statistical properties of the distribution of activations. For example, centered kernel alignment (Kornblith et al. 2019), orthogonal Procrustes distance (Schönemann 1966) and methods based on canonical correlation analysis (Raghu et al. 2017; Morcos, Raghu, and Bengio 2018) have been extensively used to analyze and compare representations (Smith et al. 2017; Raghu et al. 2021; Yadav et al. 2024). However, these structural similarity indices do not take the functionality of the networks into account directly. Accordingly, (Ding, Denain, and Steinhardt 2021) and (Hayne, Jung, and

Carter 2024) both show that these indices often fail in functional benchmarks. For example, they might not be sensitive to changes that strongly affect predictive performance, and they might be overly sensitive to non-functional differences.

Functional Similarity. Model stitching was introduced by (Lenc and Vedaldi 2015), who used 1×1 convolutional stitching layers to connect two “half-networks” in order to study the equivalence of their learned representations. (Csiszárík et al. 2021) proposed using model stitching as a framework to study the similarity of representations from a functional point of view. Along with (Bansal, Nakkiran, and Barak 2021a), they show that networks that are trained on the same task but under different settings can be stitched with minimal performance loss, which suggests that all successful networks achieve functionally similar representations. Model stitching has also been used to functionally compare representations between architecturally different models (McNeely-White, Beveridge, and Draper 2020; Yang et al. 2022) and between robust and non-robust networks (Jones et al. 2022; Balogh and Jelasity 2023).

Recently, the reliability of stitching for analyzing similarity has been questioned by (Hernandez et al. 2022), who showed that, in many cases, distant layers of classifiers can be stitched with high accuracy. Our work sheds light on the reason behind anomalies of this type.

Applied stitching. Many works have applied stitching as a method of combining networks, not focusing on similarity (Teerapittayanon et al. 2023; Guijt et al. 2024). (Pan, Cai, and Zhuang 2023; He et al. 2024; Pan et al. 2023) combine models of different sizes with stitching to achieve the optimal performance at various resource constraints. Our work suggests that the success of such applications is not necessarily due to some underlying similar representations.

2 Preliminaries

Here, we summarize the basic concepts related to model stitching and also introduce the similarity indices that we will use in our comparisons later on.

2.1 Functional Similarity: Model Stitching

Our notation closely follows that of (Balogh and Jelasity 2023). Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a feedforward neural network with m layers: $f = f_m \circ \dots \circ f_1$ where $f_i : \mathcal{A}_{f,i-1} \rightarrow \mathcal{A}_{f,i}$ maps the activation space of layer $i - 1$ to that of layer i . By definition, $\mathcal{A}_{f,0} = \mathcal{X}$. Since model stitching is a framework where two “half-networks” are connected, we introduce the notations $f_{\leq i} = f_i \circ \dots \circ f_1$ and $f_{> i} = f_m \circ \dots \circ f_{i+1}$.

Given two frozen networks f and g , and one layer from each network f_i and g_j , the abstract goal of stitching is to find out if $g_{> j}$, which we will refer to as the *receiver*, can achieve its function using the representation of $f_{\leq i}$, which we will call the *source*. In examining this, we attempt to find a transformation $T : \mathcal{A}_{f,i} \rightarrow \mathcal{A}_{g,j}$ such that the composite (or stitched) network $g_{> j} \circ T \circ f_{\leq i}$ is functionally similar to g .

Throughout this paper we will refer to T as the *stitching layer*, and for an input x , we will call $f_{\leq i}(x)$ its *source representation* and $g_{\leq j}(x)$ its *target representation*. Using these

terms, the goal of stitching is to find a transformation that matches the source representation to the target. In the following, we will discuss the constraints on T and our methods for finding the optimal transformation.

The complexity of T . A sensible requirement on T is to not increase the complexity of the stitched network significantly, but be expressive enough to allow for non-trivial mappings. The general consensus is that affine transformations are suitable (Csiszárík et al. 2021; Bansal, Nakkiran, and Barak 2021a) and we adopt this approach in this work.

For convolutional representations of sizes $\mathbb{R}^{c_1 \times h \times w}$ and $\mathbb{R}^{c_2 \times h \times w}$, where c_1 and c_2 are the number of feature maps, and h and w are the spatial dimensions of these feature maps, the stitching transformation is a 1×1 convolutional layer that applies the same affine mapping $M : \mathbb{R}^{c_1} \rightarrow \mathbb{R}^{c_2}$ to each of the $h \times w$ spatial positions.

If the representations have different spatial dimensions, a 2D resizing operation of each feature map precedes the stitching transformation. We used bilinear interpolation.

For transformer embeddings of sizes $\mathbb{R}^{n \times d_1}$ and $\mathbb{R}^{n \times d_2}$, where n is the number of tokens and d_1 and d_2 are the embedding dimensions, we apply the same affine mapping $M : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2}$ to every token.

Quantifying functional similarity. Following (Csiszárík et al. 2021), the functional similarity of the two selected layers will be quantified by the stitched network’s performance, relative to the performance of the target network g . In our case, the functionality of interest is accuracy over a supervised classification problem, given by an underlying distribution $p(x, y)$ over $\mathcal{X} \times \mathcal{Y}$. Thus, similarity is characterized by relative accuracy.

Task loss matching (TLM). Given that we treat functional performance as a measure of similarity, a natural way to find the optimal stitching transformation is through solving the learning task

$$\arg \min_{\theta} \mathbb{E}_{p(x,y)} [\mathcal{L}([g_{> j} \circ T_{\theta} \circ f_{\leq i}](x), y)] \quad (1)$$

for a suitable surrogate loss function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, over a dataset $D = \{(x_i, y_i)\}_{i=1}^n$ from the distribution $p(x, y)$. We only train the parameters θ of T_{θ} , while $g_{> j}$ and $f_{\leq i}$ are frozen, as mentioned before.

Direct matching (DM). Another option is to directly match the source representation to the target through solving the optimization problem

$$\arg \min_{\theta} \mathbb{E}_{p(x)} [\| [T_{\theta} \circ f_{\leq i}](x) - g_{\leq j}(x) \|_F]. \quad (2)$$

Note that this problem does not depend on y , hence it is independent of the classification task. Given that our choice of T_{θ} is a common affine transformation for each spatial position or token, sampling a set of equations

$$[T_{\theta} \circ f_{\leq i}](x_k) = g_{\leq j}(x_k), \quad x_k \sim p(x), \quad k = 1, \dots, K, \quad (3)$$

results in a linear problem in θ . We compute θ using the Moore-Penrose pseudoinverse.

2.2 Structural Similarity Indices

Structural similarity indices compare the internal activations of two layers over the same set of n examples. The goal is to quantify the similarity between two activation matrices $A \in \mathbb{R}^{n \times p_1}$ and $B \in \mathbb{R}^{n \times p_2}$, where p_1 and p_2 are the number of neurons. Let us now very briefly introduce the basic notions of three popular similarity indices that we will use in our evaluation. In the following, we assume that both A and B have centered columns.

Linear CKA. Centered kernel alignment (CKA) (Cortes, Mohri, and Rostamizadeh 2012) measures the alignment of the representations’ inter-example similarity structures, essentially computing an inner product of the vectors of pairwise similarities among the n examples in the two representations. (Kornblith et al. 2019) report similar results between the linear and RBF kernels used for measuring similarity between examples, so in this paper we focus on the linear variant of CKA, which can be computed as

$$\text{LCKA}(A, B) = \frac{\|B^T A\|_F^2}{\|A^T A\|_F \|B^T B\|_F}. \quad (4)$$

PWCCA. Canonical Correlation Analysis (CCA) has also been used to compute similarity indices (Raghu et al. 2017; Morcos, Raghu, and Bengio 2018). The canonical correlation coefficients ρ_i are given by

$$\rho_i = \max_{\mathbf{w}_A^i, \mathbf{w}_B^i} \text{corr}(A\mathbf{w}_A^i, B\mathbf{w}_B^i) \quad (5)$$

subject to the orthogonality conditions $\langle A\mathbf{w}_A^i, B\mathbf{w}_A^j \rangle = 0$ and $\langle A\mathbf{w}_B^i, B\mathbf{w}_B^j \rangle = 0$ for $j < i$.

The coefficients ρ_i can be used to compute a similarity score. The projection-weighted average of the coefficients was proposed by (Morcos, Raghu, and Bengio 2018), which was shown to be the best performing CCA variant in the benchmarks of (Ding, Denain, and Steinhardt 2021), given by

$$\text{PWCCA}(A, B) = \frac{\sum_i \alpha_i \rho_i}{\sum_i \alpha_i}, \quad \alpha_i = \sum_j |\langle A\mathbf{w}_A^i, \mathbf{a}_j \rangle| \quad (6)$$

where \mathbf{a}_j is the j^{th} column of A .

OPD. The orthogonal Procrustes problem asks for an orthogonal transformation of A such that it is closest to B in the Frobenius norm:

$$R^* = \arg \min_R \|B - AR\|_F^2 \quad \text{subject to} \quad R^T R = I. \quad (7)$$

The minimum distance (that we get when using the optimal transformation R^*), also known as the Orthogonal Procrustes Distance (OPD), can be computed indirectly without computing R^* as

$$\text{OPD}(A, B) = \|A\|_F^2 + \|B\|_F^2 - 2\|B^T A\|_* \quad (8)$$

where $\|\cdot\|_*$ denotes the nuclear norm. (Ding, Denain, and Steinhardt 2021) show that OPD consistently outperforms PWCCA and LCKA in their benchmarks. Note that OPD is a distance function, not a similarity score, and we use it accordingly.

2.3 Two-Faced Similarity Indices

Some of the methods can be used both as a functional and a structural similarity index. Direct matching, for example, creates an approximation of the target representation using the source representation, where the Frobenius norm of the approximation error can serve as a structural distance metric. Similarly, the optimal orthogonal transformation R^* in eq. (7) can also be interpreted as a variant of direct matching, using AR^* as the input to the second half-network. (Csiszárík et al. 2021) have argued that this functional variant of OPD is not a promising approach. However, structural direct matching is an interesting baseline that we will include in our evaluations.

3 The Unreliability of Task Loss Matching

Here, we will demonstrate that task loss matching fails the following two important sanity checks rather spectacularly:

Inter-network layer identification. (Kornblith et al. 2019) argue that between two architecturally identical networks trained from different initializations, for every layer in one network, *the architecturally corresponding layer should be the most similar one* in the other network. They showed that, by far, LCKA performs best in this test.

Intra-network layer identification. We verify the seemingly obvious requirement that within one network, for every layer the most similar layer should be itself.

Self-stitching. Previous studies, such as (Csiszárík et al. 2021; Bansal, Nakkiran, and Barak 2021b; Balogh and Jelasity 2023) focused on model stitching between the corresponding layers of architecturally identical networks. However, the concept of stitching is not restricted to corresponding layers or different networks. Here, to test the intra-network sanity check, we will use *self-stitching* where the stitching layer connects layers from the same network. This essentially means cutting out layers, or replicating layers within a network if the two stitched layers are not the same. If they are the same, self-stitching simply inserts a stitching transformation after the given layer.

3.1 Experimental Setup

We conduct experiments with two types of image classifiers: ResNets (He et al. 2016) and Vision Transformers (ViTs) (Dosovitskiy et al. 2021). We used the ResNet-18 and ViT-Ti architectures in our experiments. We trained 10 networks of each type on the CIFAR-10 (Krizhevsky 2009) dataset with identical hyperparameters, but different random initializations.

The layers we considered were all the residual blocks and the transformer encoder blocks of our ResNets and ViTs, respectively. All stitchers were initialized using the direct matching method described in section 2.1 with $K = 100$ samples from the CIFAR-10 training set. Task loss matching was then trained for 30 epochs with identical parameters as in (Csiszárík et al. 2021). All the similarity indices were evaluated over the CIFAR-10 test set. For further training details, please see the Supplementary.

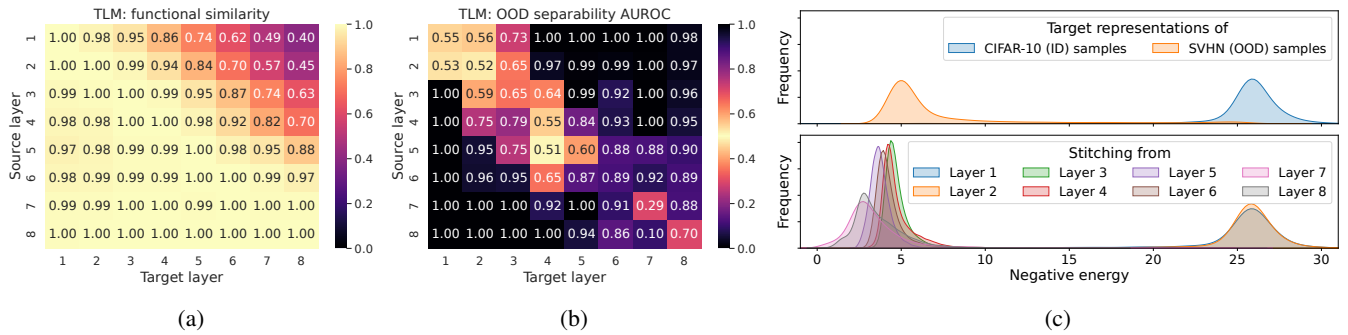


Figure 1: ResNet-18, CIFAR-10, intra-network results with task loss matching. (a) Pairwise similarities of layers (relative accuracy). (b) OOD level of the target representations (AUROC). (c) Energy distributions used by the OOD detector, when layer 1 is the target layer. (c) top: energy distributions of layer 1 on in-distribution and OOD samples; (c) bottom: energy distributions created by stitching from every layer, when measured over in-distribution samples. These distributions correspond to column 1 of the OOD matrix in (b).

Similarity index	Intra-network		Inter-network	
	RN-18	ViT-Ti	RN-18	ViT-Ti
PWCCA	100%	100%	12.50%	8.33%
OPD	100%	100%	19.17%	18.33%
LCKA	100%	100%	96.11%	34.81%
TLM	63.75%	24.17%	35.28%	10.00%
DM (struct.)	100%	100%	92.50%	25.19%
DM (func.)	100%	100%	66.39%	11.85%

Table 1: Accuracies of identifying the same layer within a network and the architecturally corresponding layers between networks based on maximum similarity.

3.2 Results

We measured intra-network similarities for all networks, and inter-network similarities between all pairs of networks. Table 1 contains the accuracies of identifying the architecturally corresponding layers based on maximum similarity.

Failed inter-network sanity check. We can see that task loss matching performs very poorly at identifying the corresponding layers between different networks. However, PWCCA and OPD also struggle with this, while LCKA performs very well, confirming the results in (Kornblith et al. 2019). Interestingly, structural direct matching (see section 2.3) achieves comparable results to LCKA and its functional version outperforms task loss matching as well. Also, all methods struggle on the ViT architecture, giving further evidence to the findings of (Raghu et al. 2021).

Failed intra-network sanity check. More alarmingly, task loss matching often fails to indicate that a layer is most similar to itself within a network. This is quite problematic because passing this sanity check can be considered a bare minimum requirement for any definition of similarity.

This failure is due to the training process that creates the stitching transformation, because the training is always initialized with the help of direct matching that does not fail this test. Let us now explore this problem by looking at the similarity values obtained with task loss matching.

3.3 A Closer Look at Intra-Network Similarities

Figure 1a shows a *similarity matrix* corresponding to our ResNet-18 experiments. The matrix shows the functional similarity according to task loss matching between every layer pair, with the source and the target layers indicated on the vertical and horizontal axes, respectively.

Clearly, task loss matching indicates high similarity values between layers that are very far apart, as was also shown in (Hernandez et al. 2022). More strikingly, the similarity between very far layers can be just as high as the similarity between the corresponding layer (which, in the case of self-stitching, is the same layer).

This is unexpected, as neural networks are thought to develop layers with different functions: earlier layers tend to be generic, while the final layers conform to the task at hand (Lenc and Vedaldi 2019). In the case of classifiers, the final layers tend to cluster representations in the activation space (Goldfeld et al. 2019; Pappayan, Han, and Donoho 2020; Balogh and Jelasity 2023; Yang, Steinhardt, and Hu 2023).

A possible explanation. To understand why task loss matching indicates a high similarity between these functionally different representations, we hypothesize that directly optimizing the stitcher for functional performance causes the stitching layer to create *out-of-distribution (OOD)* representations that fool the receiver to achieve good accuracy. This would be somewhat surprising because of the relatively weak affine transformations that form the stitching layer. To verify our hypothesis, we perform an OOD analysis of stitching.

4 Out-of-Distribution Representations

In order to study whether the internal representations are out-of-distribution (OOD), first, we need to select a suitable method for OOD detection and adapt it to our setting.

Energy-based OOD detection. For classification tasks, energy-based OOD detection (Grathwohl et al. 2020; Liu et al. 2020) has proven to be a successful approach. This method assigns an energy score—namely the negative

LogSumExp(\cdot) of the logits of the classifier—to each input sample and then finds an energy threshold that best separates the energy scores of in-distribution (ID) and OOD samples. In our experiments, we adopt the framework of (Liu et al. 2020) who pre-train a classifier on ID samples and then fine-tune it with an auxiliary OOD dataset to separate the energy scores of ID and OOD samples.

A dedicated OOD detector for each layer. We are interested in the distribution of the activations of a given layer, as opposed to that of the input samples. To this end, we train a dedicated OOD detector network for every layer of interest, using the energy-based method of (Liu et al. 2020). The training dataset of this detector network for layer i consists of the target activations $\{g_{\leq i}(x_k)\}_{k=1}^K$, where $x_k \sim p(x)$, $k = 1, \dots, K$. In this dataset, the label of the activation $g_{\leq i}(x_k)$ is the same as that of x_k . This dataset is then considered in-distribution. Similarly, the auxiliary dataset can be generated using the target activations of an OOD dataset.

4.1 Experimental Setup

We perform our experiments using the CIFAR-10 dataset, where OOD detection is known to work well. We use the 300K Random Images dataset (Hendrycks, Mazeika, and Dietterich 2019) as the auxiliary OOD dataset for fine-tuning. As OOD detectors, we use ResNet-18 models for convolutional representations and ViT-Ti models without the initial embedding sections for vision transformer representations. We trained OOD detectors on the activations of all stitched layers of every stitched network separately. Additional details are discussed in the Supplementary.

4.2 Results

After training the OOD detectors, we can apply them to test if the distribution of the stitched activations $\{[T_\theta \circ f_{\leq i}](x_k)\}_{k=1}^K$ can be separated from the target activations, and thus detected as OOD. We report the separability of the target and the stitched representations by the OOD detector with the help of the area under the receiver operating characteristic curve (AUROC) metric. The OOD AUROC values for each layer pair are presented as a heatmap for easy comparison with the similar accuracy heatmaps (for example, fig. 1b). Figure 2 present the results for the experimental settings we evaluated.

TLM prefers OOD representations. The results confirm our earlier hypothesis, namely that TLM results in internal representations that are predominantly OOD. In fact, in most cases a complete separation is possible (see fig. 1c, and the many AUROC values close to 1 in fig. 2). The only exceptions are the cases where we match an early layer to itself (intra-network), or to the architecturally corresponding layer (inter-network). At the same time, the relative accuracy of the stitched networks is predominantly excellent, even when stitching distant layers.

What is even more striking is that TLM often creates *OOD representations for architecturally matching layers* as well for the ViT architecture and also for ResNets towards the end of the network, as shown in both fig. 1b and fig. 2.

In the intra-network setting, even some identical layers are matched with OOD representations, *despite being initialized with the identity mapping*.

DM prefers ID representations. The objective of direct matching given in eq. (2) promotes creating in-distribution representations. Indeed, in fig. 2 we can observe that, when direct matching indicates high functional similarity, it does so by creating ID representations. Conversely, when the mapped representation is OOD, direct matching indicates low similarity.

OOD is undesirable. The results discussed here support the hypothesis that the reason behind TLM failing the sanity checks in section 3 is the emergence of OOD representations. Since TLM can be considered a fine-tuning of DM transformations (since TLM learning is initialized with DM), it is clear that TLM can modify the DM transformation significantly to gain maximum task performance, leading to the violation of very natural sanity checks. The OOD representations are therefore undesirable in a stitching-based functional similarity approach. It should be noted, though, that TLM is still a suitable method to provide evidence for the *lack of functional similarity*.

5 Statistical Tests for Functional Similarity

Here, we compare direct matching to structural similarity indices using the methodology of (Ding, Denain, and Steinhart 2021). Ding et al. focus on the functional aspects of similarity. Their main requirement is that *a good notion of similarity should be correlated to functional similarity*.

They define a notion of functional similarity with the help of *probing accuracy* (Belinkov et al. 2017), where a simple probing network is trained on a given representation. If two representations achieve different accuracies, they are considered functionally different. They then apply rank-correlation statistics to test whether a given similarity index is correlated with functional similarity in a number of scenarios.

In the following, we describe the basic intuition behind the testing framework and the scenarios. Further details are given in the Supplementary.

Sensitivity test. Here, the intuition is that low-rank approximations of a representation differ functionally from the original representation (based on probing accuracy), so similarity indices should also pick up on such differences. That is, they should be *sensitive to functional differences*. This can be tested via generating a series of low-rank approximations of different ranks for a fixed layer, and computing the correlation of functional (probing accuracy) and structural (similarity index) differences.

Specificity test. Here, like in our previous sanity checks, the intuition is that architecturally corresponding layers are functionally very close (based on probing accuracy) so similarity indices should also indicate the highest similarity with the architecturally corresponding layers, no matter which network they come from. That is, they should be *specific to functional differences*, and they should ignore non-functional differences such as random initialization. This




Figure 2: Intra-network and inter-network stitching similarities (relative accuracy) and the corresponding OOD scores (AUROC) with task loss matching (TLM) and direct matching (DM). OOD scores close to 0.5 indicate in-distribution. The horizontal and vertical index is the target layer and the source layer, respectively. The stitched models were the same in each row.

can be tested via generating several instances of the representations of every layer of a given architecture (training with different random initializations) and again, computing the correlations of functional and structural differences.

The testing framework. More formally, the testing framework of (Ding, Denain, and Steinhardt 2021) requires a set of representations S and the probing accuracy F defined over representations. They select a reference representation

$$A = \arg \max_{a \in S} F(a), \quad (9)$$

and for every representation $B \in S$ they compute $|F(A) - F(B)|$ and $d(A, B)$, where d is a measure of dissimilarity (eg. OPD, 1-CKA, etc.). Finally, they report the rank correlation between $|F(A) - F(B)|$ and $d(A, B)$, specifically, Kendall’s τ and Spearman’s ρ .

5.1 Experimental Setup

We present the key design choices here, the full set of details are discussed in the Supplementary.

Models. We study ResNet-18 and ViT-Ti models trained on the CIFAR-10, SVHN (Netzer et al. 2011) and ImageNet-1k (Russakovsky et al. 2015) datasets. We trained 10 instances of both architectures on CIFAR-10 and SVHN and 5 instances of both architectures on ImageNet. The training used the same hyperparameters within each scenario, but different initializations for the model instances.

Test settings. In the specificity test, we analyze all the 8 ResNet blocks and all the 12 ViT blocks.

In the sensitivity test, we include the last 4 ResNet blocks using low rank approximations of 13 and 14 different ranks for blocks 5 and 6, and blocks 7 and 8, respectively. Similarly, we studied the last 6 ViT blocks for CIFAR-10 and SVHN and the last 4 blocks for ImageNet at 16 different ranks. The ranks ranged from full-rank to rank 1.

Task	Model	LCKA		PWCCA		OPD		DM (func.)		DM (struct.)	
		τ	ρ	τ	ρ	τ	ρ	τ	ρ	τ	ρ
CIFAR-10	ResNet	0.5687	0.7217	0.6276	0.7767	0.5484	0.7176	0.7896	0.9169	0.6583	0.8068
	ViT	0.4851	0.6320	0.8260	0.9453	0.7321	0.8621	0.7448	0.8885	0.7856	0.9057
SVHN	ResNet	0.6924	0.8331	0.6913	0.8280	0.5740	0.7344	0.7617	0.8932	0.7051	0.8358
	ViT	0.7039	0.8379	0.6179	0.7823	0.6967	0.8460	0.6718	0.8636	0.7191	0.8789
ImageNet	ResNet	0.7762	0.8998	0.8315	0.9238	0.6696	0.8456	0.8096	0.9141	0.6631	0.7507
	ViT	0.5945	0.7518	0.8355	0.9458	0.6421	0.8146	0.7757	0.9209	0.7737	0.9025

Table 2: Kendall’s τ and Spearman’s ρ rank correlations for the sensitivity test.

Task	Model	LCKA		PWCCA		OPD		DM (func.)		DM (struct.)	
		τ	ρ	τ	ρ	τ	ρ	τ	ρ	τ	ρ
CIFAR-10	ResNet	0.8590	0.9699	0.7285	0.8959	0.7678	0.9202	0.8571	0.9698	0.8760	0.9749
	ViT	0.7189	0.8701	0.6502	0.7848	0.5314	0.6818	0.7818	0.9207	0.6639	0.8019
SVHN	ResNet	0.9288	0.9916	0.7610	0.9136	0.9269	0.9916	0.8895	0.9819	0.9006	0.9845
	ViT	0.7148	0.8755	0.3015	0.4245	≈ 0	≈ 0	0.7415	0.8931	0.2898	0.4095
ImageNet	ResNet	0.8436	0.9516	0.6590	0.8214	0.5974	0.6784	0.8564	0.9643	≈ 0	≈ 0
	ViT	0.7141	0.8645	0.4508	0.5965	0.4994	0.6356	0.7912	0.9179	0.4911	0.6443

Table 3: Kendall’s τ and Spearman’s ρ rank correlations for the specificity test. The value ≈ 0 indicates that the given correlation test is not statistically significant.

In the analysis of direct matching, we applied SVD-based low rank approximation (LRA) to the source and target representations before computing the parameters of the optimal transformation. We also used LRA on the source representations during evaluation.

Since our work is focused on image classifiers, we chose F to be the linear probing accuracy of a representation (Alain and Bengio 2017).

5.2 Results

Table 2 and Table 3 show the results of the sensitivity and the specificity test, respectively. We show the mean of the layer-wise results of the sensitivity test, as the results are consistent for individual layers.

Direct matching performs well. Functional direct matching achieves consistently high rank correlations with the functionality of interest in both tests. In fact, it often has the highest rank correlations among all similarity measures, while being close to the leader in most of the other cases. Therefore, we argue that direct matching can indeed serve as a meaningful measure of similarity, despite being a structural similarity at its core. Based on direct matching’s previously discussed tendency to create in-distribution representations, and its good performance in structural and functional benchmarks, direct matching is a promising baseline for both structural and functional similarity.

The conclusions of Ding et al. revisited. As concluded by (Ding, Denain, and Steinhardt 2021), LCKA performs poorly in the sensitivity test, while PWCCA achieves low rank correlations in the specificity test. However, OPD—the similarity index recommended by Ding et al.—shows inconsistent results, which questions the conclusion that OPD would be a preferable baseline for functional evaluations.

6 Conclusions and Limitations

We argued that, although functional similarity is an important concept, focusing purely on function can result in behavior that is strongly inconsistent with basic sanity checks for similarity indices. In fact, extending on the observation of (Hernandez et al. 2022), we even show that TLM-based model stitching (the method of choice for functional similarity) can indicate that, for a given layer, there are more similar layers than the layer itself, which is highly counter-intuitive.

We demonstrated that the reason behind these problems is that TLM-based stitching generates highly OOD representations. In fact, TLM can result in highly OOD representations even if architecturally identical layers are stitched.

Based on our empirical evaluation, we concluded that DM-based stitching combines structural and functional perspectives and as such it generates in-distribution representations when matching is possible, while being able to assess functional compatibility as well.

It has to be mentioned that our evaluation methodology is rather expensive. Table 1 alone requires computing 22,880 stitching transformations. Altogether, we computed around 40,000 transformations, trained 60 models, 130 OOD detectors, and 3590 probing layers. Counting also our preliminary experiments, we used more than 1.5 GPU year’s worth of computation on a mixture of GPUs from RTX 2080TI to RTX A6000.

As for limitations, it would be useful to verify our claims on a wider set of models and applications, including ones outside the image processing domain. While in principle it is enough to demonstrate one problematic scenario to prove that a method is not reliable, it is not clear how generic the discovered problems are and how they depend on task-complexity and model complexity. This, however, would require orders of magnitudes more resources than what is available to us.

Acknowledgements

This work was supported by the University Research Grant Program of the Ministry for Culture and Innovation from the source of the National Research, Development and Innovation Fund, the European Union project RRF-2.3.1-21-2022-00004 within the framework of the Artificial Intelligence National Laboratory, and by the project TKP2021-NVA-09, implemented with the support provided by the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation Fund, financed under the TKP2021-NVA funding scheme.

References

- Alain, G.; and Bengio, Y. 2017. Understanding intermediate layers using linear classifier probes. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Balogh, A.; and Jelasity, M. 2023. On the Functional Similarity of Robust and Non-Robust Neural Representations. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 1614–1635. PMLR.
- Bansal, Y.; Nakkiran, P.; and Barak, B. 2021a. Revisiting model stitching to compare neural representations. *Advances in neural information processing systems*, 34: 225–236.
- Bansal, Y.; Nakkiran, P.; and Barak, B. 2021b. Revisiting Model Stitching to Compare Neural Representations. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 225–236.
- Belinkov, Y.; Durrani, N.; Dalvi, F.; Sajjad, H.; and Glass, J. R. 2017. What do Neural Machine Translation Models Learn about Morphology? In Barzilay, R.; and Kan, M., eds., *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, 861–872. Association for Computational Linguistics.
- Cortes, C.; Mohri, M.; and Rostamizadeh, A. 2012. Algorithms for Learning Kernels Based on Centered Alignment. *J. Mach. Learn. Res.*, 13: 795–828.
- Csiszárík, A.; Kőrösi-Szabó, P.; Matszangosz, A.; Papp, G.; and Varga, D. 2021. Similarity and matching of neural network representations. *Advances in Neural Information Processing Systems*, 34: 5656–5668.
- Csiszárík, A.; Kőrösi-Szabó, P.; Matszangosz, Á. K.; Papp, G.; and Varga, D. 2021. Similarity and Matching of Neural Network Representations. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 5656–5668.
- Ding, F.; Denain, J.-S.; and Steinhardt, J. 2021. Grounding representation similarity through statistical testing. *Advances in Neural Information Processing Systems*, 34: 1556–1568.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Goldfeld, Z.; van den Berg, E.; Greenewald, K. H.; Melnyk, I.; Nguyen, N.; Kingsbury, B.; and Polyanskiy, Y. 2019. Estimating Information Flow in Deep Neural Networks. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 2299–2308. PMLR.
- Grathwohl, W.; Wang, K.-C.; Jacobsen, J.-H.; Duvenaud, D.; Norouzi, M.; and Swersky, K. 2020. Your classifier is secretly an energy based model and you should treat it like one. In *International Conference on Learning Representations*.
- Guijt, A.; Thierens, D.; Alderliesten, T.; and Bosman, P. A. 2024. Stitching for Neuroevolution: Recombining Deep Neural Networks without Breaking Them. *arXiv preprint arXiv:2403.14224*.
- Hayne, L.; Jung, H.; and Carter, R. 2024. Does Representation Similarity Capture Function Similarity? *Transactions on Machine Learning Research*.
- He, H.; Pan, Z.; Liu, J.; Cai, J.; and Zhuang, B. 2024. Efficient Stitchable Task Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28555–28565.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hendrycks, D.; Mazeika, M.; and Dietterich, T. 2019. Deep Anomaly Detection with Outlier Exposure. *Proceedings of the International Conference on Learning Representations*.
- Hernandez, A.; Dangovski, R.; Lu, P. Y.; and Soljagic, M. 2022. Model Stitching: Looking For Functional Similarity Between Representations. In *SVRHM 2022 Workshop @ NeurIPS*.
- Jones, H.; Springer, J. M.; Kenyon, G. T.; and Moore, J. 2022. If You’ve Trained One You’ve Trained Them All: Inter-Architecture Similarity Increases With Robustness. In *The 38th Conference on Uncertainty in Artificial Intelligence*.
- Kornblith, S.; Norouzi, M.; Lee, H.; and Hinton, G. 2019. Similarity of neural network representations revisited. In *International conference on machine learning*, 3519–3529. PMLR.

- Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto.
- Lenc, K.; and Vedaldi, A. 2015. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 991–999.
- Lenc, K.; and Vedaldi, A. 2019. Understanding Image Representations by Measuring Their Equivariance and Equivalence. *International Journal of Computer Vision*, 127(5): 456–476.
- Liu, W.; Wang, X.; Owens, J. D.; and Li, Y. 2020. Energy-based Out-of-distribution Detection. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- McNeely-White, D.; Beveridge, J. R.; and Draper, B. A. 2020. Inception and ResNet features are (almost) equivalent. *Cognitive Systems Research*, 59: 312–318.
- Morcos, A.; Raghu, M.; and Bengio, S. 2018. Insights on representational similarity in neural networks with canonical correlation. *Advances in neural information processing systems*, 31.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, 4. Granada.
- Pan, Z.; Cai, J.; and Zhuang, B. 2023. Stitchable Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16102–16112.
- Pan, Z.; Liu, J.; He, H.; Cai, J.; and Zhuang, B. 2023. Stitched ViTs are Flexible Vision Backbones. *arXiv preprint arXiv:2307.00154*.
- Papayan, V.; Han, X.; and Donoho, D. L. 2020. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40): 24652–24663.
- Raghu, M.; Gilmer, J.; Yosinski, J.; and Sohl-Dickstein, J. 2017. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. *Advances in neural information processing systems*, 30.
- Raghu, M.; Unterthiner, T.; Kornblith, S.; Zhang, C.; and Dosovitskiy, A. 2021. Do Vision Transformers See Like Convolutional Neural Networks? In Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3): 211–252.
- Schönemann, P. H. 1966. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1): 1–10.
- Smith, S. L.; Turban, D. H. P.; Hamblin, S.; and Hammerla, N. Y. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. In *International Conference on Learning Representations*.
- Teerapittayanon, S.; Comiter, M.; McDanel, B.; and Kung, H. 2023. StitchNet: Composing Neural Networks from Pre-Trained Fragments. In *2023 International Conference on Machine Learning and Applications (ICMLA)*, 61–68.
- Yadav, S.; Theodoridis, S.; Hansen, L. K.; and Tan, Z.-H. 2024. Masked Autoencoders with Multi-Window Local-Global Attention Are Better Audio Learners. In *The Twelfth International Conference on Learning Representations*.
- Yang, X.; Zhou, D.; Liu, S.; Ye, J.; and Wang, X. 2022. Deep model reassembly. *Advances in neural information processing systems*, 35: 25739–25753.
- Yang, Y.; Steinhart, J.; and Hu, W. 2023. Are Neurons Actually Collapsed? On the Fine-Grained Structure in Neural Representations. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 39453–39487. PMLR.