

All Roads Lead to Rome? Exploring the Invariance of Transformers' Representations

Working Paper**Author(s):**

Ren, Yuxin; Guo, Qipeng; Jin, Zhijing; Ravfogel, Shauli; Sachan, Mrinmaya; Schölkopf, Bernhard; Cotterell, Ryan

Publication date:

2023-05-23

Permanent link:

<https://doi.org/https://doi.org/10.3929/ethz-b-000653513>

Rights / license:

[Creative Commons Attribution-ShareAlike 4.0 International](#)

Originally published in:

arXiv, <https://doi.org/10.48550/ARXIV.2305.14555>

Funding acknowledgement:

- Neuro-cognitive Model Inspired from Human Language Processing ()

All Roads Lead to Rome? Exploring the Invariance of Transformers' Representations

Yuxin Ren^{1,*} Qipeng Guo² Zhijing Jin^{3,4} Shauli Ravfogel⁵
Mrinmaya Sachan⁴ Bernhard Schölkopf^{3,4} and Ryan Cotterell⁴

¹Tsinghua University ²Fudan University

³Max Planck Institute for Intelligent Systems, Tübingen, Germany

⁴ETH Zürich ⁵Bar-Ilan University

jinzhi@ethz.ch ryx20@mails.tsinghua.edu.cn

Abstract

Transformer models bring propelling advances in various NLP tasks, thus inducing lots of interpretability research on the learned representations of the models. However, we raise a fundamental question regarding the reliability of the representations. Specifically, we investigate whether transformers learn essentially isomorphic representation spaces, or those that are sensitive to the random seeds in their pretraining process. In this work, we formulate the *Bijection Hypothesis*, which suggests the use of bijective methods to align different models' representation spaces. We propose a model based on invertible neural networks, BERT-INN, to learn the bijection more effectively than other existing bijective methods such as the canonical correlation analysis (CCA). We show the advantage of BERT-INN both theoretically and through extensive experiments, and apply it to align the reproduced BERT embeddings to draw insights that are meaningful to the interpretability research.¹

1 Introduction

Transformer models (Vaswani et al., 2017) have emerged as a powerful architecture for natural language processing (NLP) tasks, achieving state-of-the-art results on a wide range of benchmarks (Devlin et al., 2019; Brown et al., 2020; OpenAI, 2023, *inter alia*). However, one issue that has received less attention is the robustness of these models to random initializations. Specifically, it is unclear whether Transformer models learn essentially *isomorphic representation spaces*, or if they learn representations that are sensitive to the random seeds in their pretraining process.

* Done during the research internship with Max Planck Institute & ETH Zürich.

¹Our code is at <https://github.com/twinkle0331/BERT-similarity>.

Recent studies has shown that small variations in the random seed used for training can result in significant performance differences between models with the same architecture (Sellam et al., 2022; Dodge et al., 2020). Understanding the degree to which the learned representation space is sensitive to random initialization is important because it has implications for interpretability research, i.e., understanding *what is encoded* in those models. Such efforts can include an analysis of the learned representations of LLMs by probing the information contained within these representations (Tenney et al., 2019; Rogers et al., 2020), or by observing the change of neuron activations under intervention (Ravfogel et al., 2021; Meng et al., 2022; Lasri et al., 2022). As we illustrate in Figure 1, if a model's learned representation space is highly sensitive to initialization, then it may be necessary to run probing tasks multiple times with different random initializations to obtain more reliable results. Alternatively, techniques to improve the robustness of the learned representation space can be applied to make the models more suitable for probing tasks.

The representation spaces of models trained with different random initializations may be isomorphic due to two main reasons: the properties of the training data or inherent symmetries in the models. For instance, permutation symmetry, which permits the swapping of neurons in the hidden layer without altering the loss function's value, is a typical symmetry observed in neural networks (Hecht-Nielsen, 1990). However, prior research has suggested that the permutation symmetry in neural networks may not entirely capture the underlying invariances in their training dynamics (Ainsworth et al., 2022).

This paper introduces the *Bijection Hypothesis*, which posits the existence of a bijection mapping between the representation spaces of different models. To investigate this hypothesis, we propose

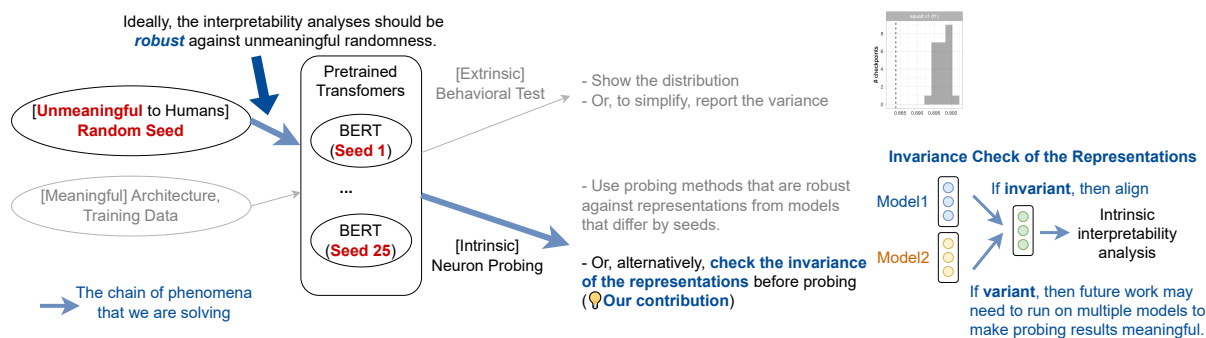


Figure 1: Overview of this work. Specifically, we investigate the similarity of reproduced BERT representations. Given multiple BERT models with different random seeds, we align their representation space according to our *Bijection Hypothesis*, and show clusters of BERT embeddings.

BERT-INN, an invertible neural network (INN) to align the representation spaces of two BERT models by a set of affine coupling layers. We show the advantage of our BERT-INN over four other bijective methods both theoretically and through three different experiments.

Finally, we apply our BERT-INN method to align the embedding spaces of the reproduced BERTs (Sellam et al., 2022), and draw several key findings. Specifically, we identify that (1) the shallower layers are, in general, more invariant than the deeper layers; (2) the interactions of neurons (i.e., attention weights) are more invariant than the representations; and (3) finetuning breaks the variance.

In conclusion, our contributions are as follows:

1. We are, to the best of our knowledge, the first to highlight the consistency check of different transformer models' representation spaces.
2. We pose the *Bijection Hypothesis*, for which we train a BERT-INN model to learn the bijective function mapping and recover the non-linear transformation.
3. We show the effectiveness our BERT-INN model over four other bijective methods in terms of both theoretical properties and three different empirical experiments.
4. We align the BERT representations using our BERT-INN method, and reveal several insights that have meaningful implications for future work on interpretability research of transformer models.

2 Background of LLM Analysis

2.1 Large Language Models (LLMs)

The recent success of NLP is driven by LLMs (Radford et al., 2018; Devlin et al., 2019; Liu et al., 2019; Brown et al., 2020; Ouyang et al., 2022). These models are pre-trained on massive amount of text data, allowing them to learn general language representations that can be fine-tuned for specific NLP tasks. Scaling up the size of language models has been shown to confer a range of benefits, such as improved performance and sample efficiency. LLMs show strong performance in few-shot or even zero-shot tasks and a better generalization of out of distribution tasks.

2.2 Interpretability of LLMs

The empirical success of BERT on various NLP domains, encourages researchers to analyse and understand these architectures (Rogers et al., 2020). Kovaleva et al. (2019) visualize self-attention modules to probe how tokens in a sequence interact with each other. Clark et al. (2019) provide in-depth analysis on the self-attention mechanism, showing that some attention heads correspond to linguistic notions of syntax and coreference. In addition, Zhao et al. (2020); Merchant et al. (2020) investigate the impacts of downstream task finetuning for BERT representations. Hewitt and Liang (2019); Pimentel et al. (2020) stress that controlling for important factors, e.g., probing classifier's capacity, is crucial to obtain meaningful probing results.

2.3 The Effect of Random Seeds

Sellam et al. (2022) conduct experiments on pre-trained models with different random seeds for

pretraining, and find that the instance-level agreement on downstream tasks varies by the random seeds, and the gap on out-of-distribution tasks is more pronounced. Dodge et al. (2020) shows that varying only the random seed of fine-tuning process leads to substantial performance gap. The two factors influenced by random seed, training order and weight initialization, contribute equally to the variance of performance. Zhong et al. (2021) argue that model predictions are noisy at the instance level. On the MNLI dataset, even the same architecture with different finetuning seeds will lead to different predictions on $\approx 8\%$ of the instances, due to under-specification (D’Amour et al., 2020).

3 The *Bijection Hypothesis*

3.1 Representation Spaces of BERTs

Given two BERT models \mathcal{M}_1 and \mathcal{M}_2 with identical architecture, but different random initializations, the goal is to evaluate their similarity through their representation spaces.

We denote the sets of BERT parameters after the pre-training as θ_1 and θ_2 , respectively. Their representation transformation functions are thus f_{θ_1} and f_{θ_2} , respectively. Since a BERT has 12 transformer layers (Vaswani et al., 2017), it is a hyperparameter to specify the representation space of which transformer layer we want to analyze. Let us denote the layer of interest as l . Then, the sequence transformation function works as follows. For any text sequence t with k tokens, we first get its token-wise embedding $\mathbf{T} \in \mathbb{R}^{k \times 768}$ at Layer l by $\mathbf{T} = \text{Transformer}_{1,\dots,l}(t) = \text{FFN}_l(\text{Att}_l(\dots \text{FFN}_1(\text{Att}_1(\text{Emb}(t)))))$, where FFN_i and Att_i denotes the feedforward layer and the self-attention network at Layer i , respectively, and Emb denotes the embedding layer. Then we follow the practice in Del and Fishel (2021); Wu et al. (2019); Zhelezniak et al. (2019) to do a mean-pooling operation to obtain the overall sentence embedding $\frac{1}{k} \sum_{i=1}^k \mathbf{T}_{i,*}$.

For the two transformation functions f_{θ_1} and f_{θ_2} , we feed a large list of n arbitrary text inputs, $\mathbf{D} := [t_1, t_2, \dots, t_n]$ to obtain the two sentence representation spaces:

$$\mathbf{X} := [f_{\theta_1}(t_1), \dots, f_{\theta_1}(t_n)] \text{ for } \mathcal{M}_1 \quad (1)$$

$$\mathbf{Y} := [f_{\theta_2}(t_1), \dots, f_{\theta_2}(t_n)] \text{ for } \mathcal{M}_2. \quad (2)$$

Here, $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times d}$, where d is the dimension of these embeddings, or the number of features.

3.2 Perfect Similarity = Diffeomorphism

We propose the *Bijection Hypothesis*, which states that two models are perfectly similar if and only if there exists a diffeomorphism in their representation spaces. The mapping can be discovered through techniques such as linear or non-linear projections.

Namely, for the two spaces \mathbf{X} and \mathbf{Y} , we learn an invertible function g , such that we minimize the distance metric δ between the transformed $g(\mathbf{X})$ and the original \mathbf{Y} . The similarity of the two spaces can be quantified by the distance metric, and if $\delta(g(\mathbf{X}), \mathbf{Y})$ is zero, then the two spaces are perfectly similar.

3.3 Reformulating Existing Similarity Indices into Linear Bijective Functions

We briefly introduce existing methods to align the two representation spaces before measuring their similarity.

Linear Regression. An intuitive way to learn a bijective function for aligning two representation spaces is linear regression, which can capture all linear transformations by minimizing the L2 norm. Namely, we learn a matrix \mathbf{W} which satisfies the following objective:

$$\min_{\mathbf{W}} \delta(\mathbf{X}\mathbf{W}, \mathbf{Y}) = \min_{\mathbf{W}} \|\mathbf{X}\mathbf{W} - \mathbf{Y}\|_2.$$

Canonical Correlation Analysis (CCA). CCA projects the two matrices \mathbf{X} and \mathbf{Y} into the subspaces which maximize the canonical correlation between their projections. The formulation is as follows:

$$\rho_i = \max_{\mathbf{w}_X^i, \mathbf{w}_Y^i} \text{corr}(\mathbf{X}\mathbf{w}_X^i, \mathbf{Y}\mathbf{w}_Y^i) \quad (3)$$

$$\text{for } \forall_{j < i} \mathbf{X}\mathbf{w}_X^i \perp \mathbf{X}\mathbf{w}_X^j, \quad (4)$$

$$\forall_{j < i} \mathbf{Y}\mathbf{w}_Y^i \perp \mathbf{Y}\mathbf{w}_Y^j. \quad (5)$$

Here, the objective is to maximize the correlation between all pairs of projected subspaces of \mathbf{X} and \mathbf{Y} , and the constraint is to keep the orthogonality among all the subspaces of each original space.

Since the dimension of CCA is equal to d , by concatenating the projection vectors into matrices \mathbf{W}_X and \mathbf{W}_Y , both \mathbf{X} and \mathbf{Y} are invertible matrices. The transformation matrix from \mathbf{X} to \mathbf{Y} is then equivalent to the product of \mathbf{W}_X and the inverse of \mathbf{W}_Y :

$$\mathbf{W}_X \mathbf{W}_Y^{-1}. \quad (6)$$

However, since we have to pre-determine the number of canonical correlations c , CCA suffers when the guessed c differs from the optimal one.

Noise-Robust Improvements of CCA. As CCA might suffer from outliers in the two spaces, there are two types of improvements that are more robust towards outliers or noises in the spaces.

The first variant is singular vector CCA (SVCCA), which performs singular value decomposition (SVD) to extract directions which capture most of the variance of the original space, thus removing the noisy dimensions before running the CCA.

The other method is projection-weighted CCA (PWCCA), which computes a weighted average over the subspaces learned by CCA:

$$\rho_{\text{PWCCA}} = \frac{\sum_{i=1}^c \alpha_i \rho_i}{\sum_{i=1}^c \alpha_i} \quad (7)$$

$$\alpha_i = \sum_j |\langle \mathbf{p}_i, \mathbf{x}_j \rangle| \quad (8)$$

$$\mathbf{p}_i = \mathbf{X} \mathbf{w}_X^i, \quad (9)$$

where \mathbf{x}_j is the j -th feature of the \mathbf{X} space, namely, the j -th column. The weight α_i for each correlation measure ρ_i is calculated as how similar the i -th subspace projection \mathbf{p}_i is to all the features \mathbf{x}_j of the original space \mathbf{X} .

However, SVCCA and PWCCA are still *linear* methods and unable to capture non-linear relationship between \mathbf{X} and \mathbf{Y} . Further, Csizsárík et al. (2021) shows that the low-rank transformation learned by SVD is not as effective as learning a matching under a task loss and the effectiveness of learning a matching stems from the non-linear function in a network.

3.4 Non-Linear Bijection by INNs

All the linear bijective methods introduced previously do not necessarily suit BERTs, as BERT learns non-linear feature transformation through the non-linear activation layers GELU (Hendrycks and Gimpel, 2016). Between two BERT models, there might be a complex, non-linear correspondence between their feature spaces. Hence, we propose the use of INNs to discover bijective mappings between the two BERTs.

INNs learn a neural network-based *invertible* transformation from the latent space \mathcal{U} to the observed space \mathcal{Z} . Using it as our bijective method, we learn an invertible transformation $g(\cdot)$ from \mathbf{X}

to \mathbf{Y} . Specifically, we optimize the parameters θ_g of the INN function by minimizing the L2 distance δ between the transformed $g_{\theta}(\mathbf{X})$ and \mathbf{Y} :

$$\min_{\theta_g} \delta(g_{\theta_g}(\mathbf{X}), \mathbf{Y}). \quad (10)$$

Intuitively, the invertibility constraint verifies that the mapping between the two representation spaces is *isomorphic*: there is no loss of information when mapping from \mathbf{X} to \mathbf{Y} .

For the implementation of the INN, we draw inspirations from RealNVP (Dinh et al., 2016) and adapt it into BERT-INN, which stacks L affine coupling layers to learn the bijective function g to map from the embedding of the text \mathbf{t} by the first BERT model $f_{\theta_1}(\mathbf{t})$, to that of the second BERT model $f_{\theta_2}(\mathbf{t})$.

In our BERT-INN, each affine coupling layer ℓ takes in an input embedding $e^{\ell-1}$ of dimension K , sets a smaller dimension $k < K$, and produces the output embedding e^ℓ in the following way:

$$e^0 = f_{\theta_1}(\mathbf{t}) \quad (11)$$

$$e_{1:k}^\ell = e_{1:k}^{\ell-1} \quad (12)$$

$$e_{k+1:K}^\ell = e_{k+1:K}^{\ell-1} \odot \exp(s^\ell(e_{1:k}^{\ell-1})) + t^\ell(e_{1:k}^{\ell-1}). \quad (13)$$

In each layer, it keeps the first k dimensions of the input as in Eq. (12), and transforms the remaining $K - k$ dimensions by a scaling layer s^ℓ and a translation layer t^ℓ as in Eq. (13). And the first layer takes in the embedding e^0 from the representations of the first BERT model as in Eq. (11). The main spirit of these affine coupling layers is that they form an INN simple enough to invert, but also relies on the remainder of the input vector in a non-trivial manner.

Finally, the loss objective for the last layer is as follows:

$$\mathcal{L} = \|e^L - f_{\theta_2}(\mathbf{t})\|_2, \quad (14)$$

where we optimize for the L2 distance $\|\cdot\|_2$ between the last layer’s embedding e^L and the second BERT’s representations $f_{\theta_2}(\mathbf{t})$.

4 Comparing the Bijective Methods

In this section, we show the effectiveness of our proposed INN method against all other bijective methods in terms of the theoretical properties, and empirical performance across synthetic data and real NLP datasets.

Dataset	Task	Train	Test	Domain
MNLI	Inference	393K	20K	Mixed
SST-2	Sentiment	67K	1.8K	Movie reviews
MRPC	Paraphrase	3.7K	1.7K	News

Table 1: Dataset statistics from the GLUE benchmark (Wang et al., 2019).

4.1 Experimental Setup

Reproduced BERT Models We conduct the experiments using the 25 reproduced BERT models released by Sellam et al. (2022). The BERT models vary only by random seeds, but adopt the same training procedure as in the original BERT model (Devlin et al., 2019).

Datasets To get the BERT representations, we run the 25 trained BERT models in the inference mode, and plug in a variety of datasets and tasks. Specifically, we choose three diverse tasks from the commonly used GLUE benchmark (Wang et al., 2019): the natural language inference dataset MNLI (Williams et al., 2018), the sentiment classification dataset SST-2 (Socher et al., 2013), and the paraphrase detection dataset MRPC (Dolan and Brockett, 2005). These datasets also cover a variety of domains from news, social media, and movie reviews, which are the common domains for many NLP datasets. The dataset statistics are in Table 1.

Baselines We compare against four commonly used existing bijective methods introduced in Section 3.4, namely linear regression, CCA, SVCCA, and PWCCA. Note that PWCCA sometimes generates results same as or similar to CCA, because it learns the mapping using the same principles as the CCA, and then computes a weighted average over the coefficients, which could generate similar results to CCA under some conditions.

Evaluation Setup For each method g , we evaluate their alignment performance by measuring the distance between two spaces: the ground-truth target space \mathbf{Y} , versus the learned space $g(\mathbf{X})$ transformed from the original \mathbf{X} with the alignment method $g(\cdot)$. We adopt the a widely used distance function, Euclidean distance, also known as the L2 distance, to measure the distance between the two spaces, namely $\delta(g(\mathbf{X}), \mathbf{Y}) := \|\mathbf{Y} - g(\mathbf{X})\|_2$. Since our goal is to align the two spaces, the smaller the L2 distance, the better, with the best

value being zero. For SVCCA, since it reduces the dimension of the original spaces, we compute the L2 norms between the dimension-reduced matrices. To show the effectiveness of our INN method, we randomly select eight model pairs to report the average of their pair-wise L2 distances using each of the bijective methods. And after we verify the effectiveness of our INN method, we apply it on all 25 models to draw overall insights on the similarity of the model representations.

4.2 Advantage of INN in Theoretical Properties

	Transformation Type	Non-Linear Noise-Robust	
Lin. Reg.	$\ Q_Y^T X\ _F^2 / \ X\ _F^2$	\times	\times
CCA	$\ Q_Y^T Q_X\ _F^2 / d$	\times	\times
SVCCA	$\ (U_Y T_Y)^T U_X T_X\ _F^2 / \min(\ T_X\ _F^2, \ T_Y\ _F^2)$	\times	\checkmark
PWCCA	$\sum_{i=1}^d \alpha_i \rho_i / \ \alpha\ _1$	\times	\checkmark
INN (Ours)	Any invertible function	\checkmark	\checkmark

Table 2: Theoretical properties of all the alignment methods. We focus on properties such as whether the method can handle non-linear mapping between the two spaces (“Non-Linear”), and whether the method is robust against noise in the transformation (“Allow Noise”).

We compare the theoretical properties of all the alignment methods in Table 2. As a result of the nature of the transformation types, we can see that all previous methods can only handle linear transformations, whereas our INN can model non-linear, invertible mappings. Moreover, linear regression and CCA are sensitive towards noise in the embedding space; SVCCA and PWCCA improves upon CCA by filtering out the noise; and our model is also robust against noises in that the training objective of INN is to minimize the overall loss across all data points.

4.3 Advantage of INN on Synthetic Data

We also show empirical evidence of the effectiveness of our method. For our empirical experiment, we want to understand whether INNs are capable of identifying the correspondence between the two representation spaces. We cannot do that with raw BERT representations, since we do not know how invariant the BERT representations are, and also we do not know its form of bijection, if such invariance exists. Therefore we generate synthetic data where the correspondence is known, and we measure the ability of INNs to identify it.

Specifically, we make the first space \mathbf{X} a real BERT representation space generated by Sellam

	L2 Distance $\delta(g(\mathbf{X}), \mathbf{Y}) (\downarrow)$
Linear Regression	1.0882 \pm 0.3867
CCA	0.3060 \pm 0.1948
SVCCA	0.1396 \pm 0.1246
PWCCA	0.3060 \pm 0.1948
INN (Ours)	0.0059 \pm 0.0037

Table 3: Empirical performance of all the alignment methods to recover a non-linear, invertible transformation. For L2 distance (\downarrow), the smaller, the better, and the best possible L2 distance is zero, which means the learned mapping recovers the ground-truth transformation between the two spaces.

et al. (2022). For the second subspace \mathbf{Y} , we generate it by applying a randomly initialized INN f to the first space \mathbf{X} , namely $\mathbf{Y} = f(\mathbf{X})$ as the ground truth.

Then, without any knowledge of the transformation f , we apply all methods to obtain their corresponding mapping function g given the two spaces \mathbf{X} and \mathbf{Y} . For each method, we evaluate their alignment performance by measuring the distance between the learned $g(\mathbf{Y})$ and ground-truth \mathbf{X} using L2 distance, namely $\delta(g(\mathbf{Y}), \mathbf{X}) := \|g(\mathbf{Y}) - \mathbf{X}\|_2$. Note that since the ground truth mapping f gives an L2 distance of zero, when interpreting the results, the lower L2 distance, the better an alignment method is.

We show the performance of our method compared with all the baselines in Table 3. All baselines (linear regression, CCA, SVCCA, and PWCCA) struggle to learn a good mapping between the two spaces, with an L2 distance always larger than 0.13, which is quite large. In contrast, our INN method gives a very small L2 distance, 0.0059, which almost recovers the ground-truth transformation. The results of this experiment echoes with the theoretical properties analyzed previously in Section 4.2.

4.4 Advantage of INN on Real Data’s Representation Similarity

We now move from synthetic data to real NLP datasets. As introduced in Section 4.1, we use three diverse datasets: MNLI (Williams et al., 2018), SST-2 (Socher et al., 2013), and MRPC (Dolan and Brockett, 2005).

We report the alignment results in Table 4. For each method, we use the splits in the original data, learn the mapping using the training set, (if applicable,) tune the hyperparameters on the validation

set, and finally report the performance on the test set. For some additional references, we also report the training set performance (to check how well the method fits the data), in addition to the test set performance (as the final alignment performance).

Among all the bijective methods in Table 4, our INN method leads to the smallest L2 distance on all datasets by a clear margin. We can also see that the vanilla CCA is usually the weakest, followed by SVCCA. PWCCA and linear regression learn the bijective mapping relatively well, but still falling behind our proposed INN method.

4.5 Advantage of INN on Real Data’s Behavioral Similarity

The results presented thus far showcase the effective mapping capability of INNs between the representation spaces of two models. However, it remains unclear how applying this mapping would impact the actual *behavior* of the models. If an accurate bijection between the representation spaces of two models is established, it should be possible to *intervene* during the forward pass of one model by substituting its representation with the corresponding representation from the other model, without causing significant loss. In this section, we conduct a behavioral analysis to investigate this hypothesis.

We design the following “layer injection” experiment: For each pair of BERTs \mathcal{M}_1 and \mathcal{M}_2 finetuned on a downstream classification task (e.g., MNLI, SST2, or MRPC), we denote the embedding spaces of the last layer before the classification as \mathbf{X} for model \mathcal{M}_1 , and \mathbf{Y} for model \mathcal{M}_2 . We learn a transformation function g to map \mathbf{X} to \mathbf{Y} . , we generate the transformed embedding $g(\mathbf{X})$ and inject it into the second model \mathcal{M}_2 . With this direct replacement, we keep all the other weights in \mathcal{M}_2 unchanged, and run it in the inference mode to get the performance on all three datasets. Our experiment design carries a similar spirit to the existing work in computer vision on convolutional neural networks (Csiszárík et al., 2021).

As we can see in Table 5, our model generates the best performance across all datasets, which is quite impressive compared with the close-to-random performance of many other methods. Moreover, our performance is even better than that of the non-bijective neural networks, which might indicate that the non-bijective transformation is

	MNLI		SST2		MRPC	
	Train	Test	Train	Test	Train	Test
CCA	0.1484 \pm 0.0763	0.1534 \pm 0.0802	0.3064 \pm 0.1791	0.2535 \pm 0.1784	0.1455 \pm 0.0663	0.1502 \pm 0.0668
SVCCA	0.1216 \pm 0.1011	0.1180 \pm 0.0989	0.2399 \pm 0.2322	0.1997 \pm 0.2375	0.1033 \pm 0.1136	0.0852 \pm 0.0730
PWCCA	0.0546 \pm 0.0441	0.0529 \pm 0.0418	0.1095 \pm 0.1214	0.0877 \pm 0.1080	0.0424 \pm 0.0344	0.0391 \pm 0.0288
Lin. Reg.	0.1003 \pm 0.0799	0.1024 \pm 0.0822	0.0884 \pm 0.0533	0.0896 \pm 0.0623	0.0618 \pm 0.0392	0.0641 \pm 0.0424
INN (Ours)	0.0520 \pm 0.0388	0.0461 \pm 0.0336	0.0563 \pm 0.0315	0.0523 \pm 0.0401	0.0299 \pm 0.0165	0.0310 \pm 0.0180
<i>Sanity Check:</i>						
Non-Bij. NN	0.0360 \pm 0.0256	0.0371 \pm 0.0272	0.0479 \pm 0.0279	0.0449 \pm 0.0352	0.0247 \pm 0.0145	0.0263 \pm 0.0165

Table 4: Alignment performance on MNLI, SST2, and MRPC datasets using CCA, SVCCA, PWCCA, linear regression (Lin. Reg.), and our INN method. We report the L2 norm δ (\downarrow) after aligning the sentence space on the training set (just for a reference) and test set (as the main result). The training and test samples were randomly selected from the respective datasets, with a sample size of 1600. For sanity check, we also report the performance of a non-bijective neural network (Non-Bij. NN), which should lead to the smallest L2 norm because it has less constraint than INN, although at the expense of violating the *Bijection Hypothesis*.

	MNLI	SST2	MRPC
Random Emb.	33.12 \pm 14.63	53.97 \pm 26.83	42.49 \pm 23.89
Random Label	33.55 \pm 0.43	50.44 \pm 1.20	52.9 \pm 2.51
Lin. Reg.	37.15 \pm 24.26	54.14 \pm 31.27	45.93 \pm 28.48
CCA	30.37 \pm 9.30	45.46 \pm 12.36	51.72 \pm 18.60
PWCCA	39.70 \pm 4.57	50.24 \pm 0.90	50.00 \pm 18.38
Non-Bij. NN	29.61 \pm 15.37	64.25 \pm 19.78	53.58 \pm 25.76
INN (Ours)	75.75 \pm 8.97	74.01 \pm 15.61	86.34 \pm 1.07

Table 5: Behavioral check (by accuracy on the validation set) of various bijective mapping methods across three datasets, MNLI, SST2, and MRPC. Apart from linear regression (Lin. Reg.), CCA, PWCCA, non-bijective neural network (Non-Bij. NN), and our INN method, we also include two random baselines, one using random labels, and the other using random embeddings (Random Emb.), where we apply an untrained, randomly initialized classifier on top of the pretrained BERT encoder.

not the proper way to go. It might distort some features in the embedding space, making it incompatible with the following layers of the injected BERT. Previous work also shows a discrepancy between representational similarity and behavior similarity.

5 Findings: Variance and Invariance across BERTs

In this section, we utilize the INN alignment method on the 25 reproduced BERTs to address the following question: *How similar are these BERTs (which only differ by random seeds)?*

This question holds significance for the NLP community because both answers matter for future work directions: If these BERTs are similar, then

all the work built on one single BERT will generalize to other BERTs. And if these BERTs differ, then follow-up work on BERT should also test their findings on the other BERTs varied by random seeds, to make sure the findings are robust.

To answer this overall question, we ask three subquestions from different perspectives: (1) Across the 12 representation layers, which ones are similar, and which ones are different? (2) How does the way how hidden states interact with each other (i.e., attention weights) vary? (3) How does finetuning enlarge or reduce the differences?

5.1 Deep vs. Shallow Layers

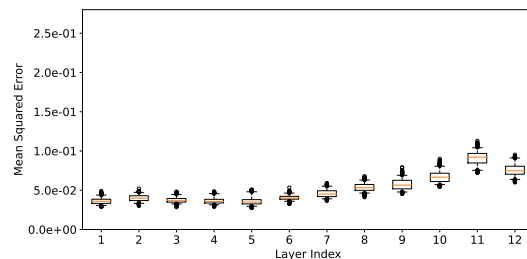


Figure 2: Similarity of the *hidden states* learned in the 12 layers of BERTs. For each layer, we align the 25 BERTs and plot the distribution of their L2 distances after alignment. The shallow layers (Layer 1 – 6) are similar, with a small L2 distance of around 0.02, whereas the deep layers (Layer 7 – 12) start to be more different, all the way to a large L2 distance of over 0.06.

We plot in Figure 2 the similarity of the representation spaces across the 12 layers of the BERTs. We show here the experimental results on the largest dataset in our previous experiments,

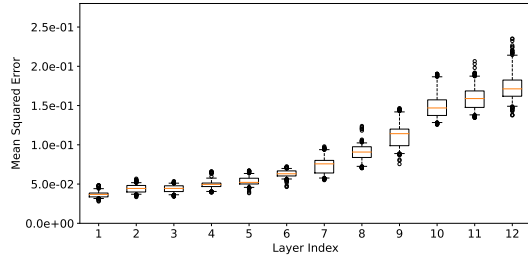


Figure 3: Similarity of the 12 layers of BERTs. We can see that the transformation in different layers are more similar for layers next to each other, and differ a lot when the layers are far away from each other.

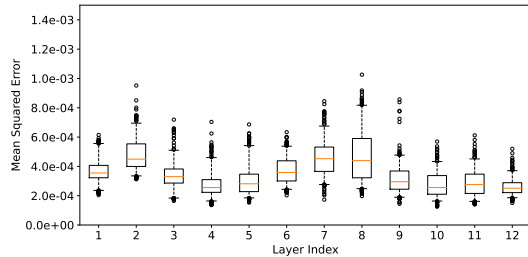


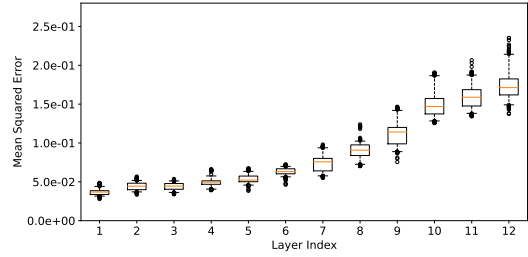
Figure 4: Similarity of the *attention weights* in the 12 layers of BERTs. For each layer, we align the 25 BERTs and plot the distribution of their L2 distances after alignment.

MNLI, and show in appendix the experiments on other datasets.

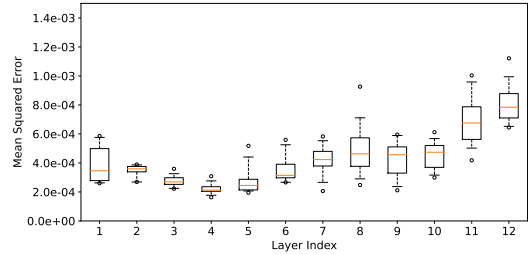
As we can see, the shallow layers (Layer 1 – 6) are relatively similar, all having an L2 distance close to 0.02, which is very small. However, the deeper layers (Layer 7 – 12) start to see an increase in inconsistency, with a larger L2 norm which rises from 0.03 to over 0.06, tripling the distance of shallow layers.

One potential explanation for the observed results is that the shallower layers of the model tend to learn more low-level features that are shared across various texts, whereas the deeper layers are capable of capturing more complex semantic features that can have multiple plausible interpretations. To support this claim, we present the results of our INN method, which measure the similarity across all 12 layers, in Figure 3.

There are three notable findings from our analysis: (1) Two areas of high similarity, represented by whiteness, are present in the plot, suggesting that the representations of the shallow layers are similar among themselves, as are the representations of the deep layers. (2) The last shows simi-



(a) Similarity of the *hidden states after finetuning*. We can see that the shallow layers stay invariant, whereas the deep layers vary even more than those before finetuning in Figure 2.



(b) Similarity of the *attention weights after finetuning*. We can see that the invariance is mostly kept across all layers.

Figure 5: Similarity across the 12 layers of BERTs after finetuning. The deep representations tend to vary more, whereas the shallow representations and attention mechanisms tend to stay similar.

ilarity to the 7th layer, since it needs both low-level features and high-level features in order to predict the word correctly. (3) The difference in the representation space changes gradually across the layers, without any abrupt changes.

5.2 Representation vs. Interaction

The transformer architecture of BERT learns not only the hidden representations, but also the way how the hidden states should *interact*, which are expressed by the weight matrices in the attention layers. Hence, we inspect the similarity of the attention weights in Figure 4. Surprisingly, the attention weights are far more invariant than the representations, with L2 distances lower by two orders of magnitude, ranging from 0.0002 to 0.0004. This indicates that the way the hidden states interact tend to be very similar across all BERTs.

5.3 Before vs. After Finetuning

Finally, we inspect how finetuning affects the similarity among BERTs. In Figure 5(a), we can see

that, although the shallow layers (1 – 6) remain similar after finetuning, the deep layers (7 – 12) vary drastically, reaching a high L2 distance of up to 0.12, which doubles its original maximum L2 distance of 0.06 in Figure 2. On the contrary, the way how the hidden states interact stays more invariant, as we can see that the attention weights stay similar across the BERTs, with L2 distances still mostly around 0.0002 to 0.0004 in Figure 5(b).

6 Conclusion

In this work, we explored the question of how invariant are transformer models that differ by random seeds. We proposed the *Bijection Hypothesis*, for which we investigated a set of bijective methods to align different BERT models. Through theoretical and empirical analyses, we identified that our proposed INN is the most effective bijective method to align BERT embeddings. Applying the INN method, we identify that shallow layers of BERTs are more consistent, whereas deeper layers encode more variance, and finetuning brings a substantial variance to the models, which imply that the finetuning process very possibly leads the model to different alternative local minima. Our work brings a unique insight for future work interpreting the internal representations of transformer models.

Ethical Considerations

This work aims to measure the similarity of the learned representation space of the commonly used NLP model, BERT. This is a step towards safer model deployment, and more transparency of the black-box LLMs. As far as we are concerned, there seem to be no obvious ethical concerns for this work. All datasets used in this work are commonly used benchmark datasets in the community, and there is no user data involved.

Author Contributions

This project originated from a discussion between Zhijing and Qipeng in 2020. Inspired by Bernhard’s causal representation learning (Schölkopf et al., 2021) and promising results in computer vision, they did some preliminary exploration into running ICA on BERT embeddings as a first step of causal representation learning for NLP. When Yuxin joins MPI and ETH as an intern with Zhijing, he discovers that ICA shows some results

that could vary a lot by the random seed, which inspires this work on inspecting the what random seeds bring to the BERT models. Yuxin then did extensive analysis into the model representations, and, then, with the insights of Ryan, we together formulate this *Bijection Hypothesis*. Yuxin completed all the analyses over 300 model pairs and many iterations of experiments. Shauli added insights when we change from CCA methods to INN, as well as many additional experiments such as exploring the behavioral similarity. Mrinmaya and Bernhard co-supervised the project. And all of the coauthors contributed to the writing.

Acknowledgment

This material is based in part upon works supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B; by the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645; by the John Templeton Foundation (grant #61156); by a Responsible AI grant by the Haslerstiftung; and an ETH Grant (ETH-19 21-1). Zhijing Jin is supported by PhD fellowships from the Future of Life Institute and Open Philanthropy.

References

- Samuel K Ainsworth, Jonathan Hayase, and Sidhartha Srinivasa. 2022. Git re-basin: Merging models modulo permutation symmetries. *arXiv preprint arXiv:2209.04836*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does

- bert look at? an analysis of bert’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286.
- Adrián Csiszárík, Péter Kőrösi-Szabó, Ákos Mat-szangosz, Gergely Papp, and Dániel Varga. 2021. Similarity and matching of neural network representations. *Advances in Neural Information Processing Systems*, 34.
- Maksym Del and Mark Fishel. 2021. Establishing interlingua in multilingual language models. *arXiv preprint arXiv:2109.01207*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Association for Computational Linguistics (ACL)*, pages 4171–4186.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. 2020. Finetuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv*.
- Bill Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Third International Workshop on Paraphrasing (IWP2005)*.
- Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. 2020. Underspecification presents challenges for credibility in modern machine learning. *Journal of Machine Learning Research*.
- Robert Hecht-Nielsen. 1990. On the algebraic structure of feedforward network weight spaces. In *Advanced Neural Computers*, pages 129–135. Elsevier.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of bert. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4365–4374.
- Karim Lasri, Tiago Pimentel, Alessandro Lenci, Thierry Poibeau, and Ryan Cotterell. 2022. Probing for the usage of grammatical number. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8818–8831.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. *arXiv preprint arXiv:2202.05262*.
- Amil Merchant, Elahe Rahimtoroghi, Ellie Pavlick, and Ian Tenney. 2020. What happens to bert embeddings during fine-tuning? *arXiv preprint arXiv:2004.14448*.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *CoRR*, abs/2203.02155.
- Tiago Pimentel, Naomi Saphra, Adina Williams, and Ryan Cotterell. 2020. [Pareto probing:](#)

- Trading off accuracy for complexity. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3138–3153, Online. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. Technical report, OpenAI.
- Shauli Ravfogel, Grusha Prasad, Tal Linzen, and Yoav Goldberg. 2021. Counterfactual interventions reveal the causal effect of relative clause representations on agreement prediction. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 194–209.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics (TACL)*, 8:842–866.
- Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. 2021. Towards causal representation learning. *CoRR*, abs/2102.11107.
- Thibault Sellam, Steve Yadlowsky, Jason Wei, Naomi Saphra, Alexander D’Amour, Tal Linzen, Jasmijn Bastings, Iulia Turc, Jacob Eisenstein, Dipanjan Das, et al. 2022. The multiberts: Bert reproductions for robustness analysis. In *ICLR*. OpenReview.net.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. Bert rediscovers the classical nlp pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Alex Wang, Amapreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations (ICLR)*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Association for Computational Linguistics (ACL)*, pages 1112–1122.
- Shijie Wu, Alexis Conneau, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Emerging cross-lingual structure in pre-trained language models. *arXiv preprint arXiv:1911.01464*.
- Mengjie Zhao, Philipp Dufter, Yadollah Yaghoobzadeh, and Hinrich Schütze. 2020. Quantifying the contextualization of word representations with semantic class probing. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1219–1234, Online. Association for Computational Linguistics.
- Vitalii Zhelezniak, April Shen, Daniel Busbridge, Aleksandar Savkov, and Nils Hammerla. 2019. Correlations between word vector sets. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 77–87.
- Ruiqi Zhong, Dhruva Ghosh, Dan Klein, and Jacob Steinhardt. 2021. Are larger pretrained language models uniformly better? comparing performance at the instance level. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3813–3827, Online. Association for Computational Linguistics.