

Introducing Orthogonal Constraint in Structural Probes

Tomasz Limisiewicz and David Mareček

Institute of Formal and Applied Linguistics, Faculty of Mathematics and Physics

Charles University, Prague, Czech Republic

{limisiewicz, marecek}@ufal.mff.cuni.cz

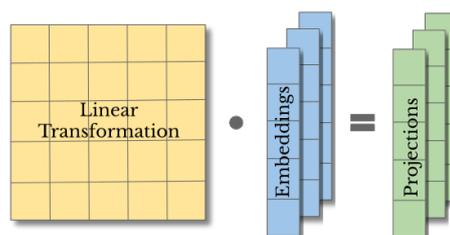
Abstract

With the recent success of pre-trained models in NLP, a significant focus was put on interpreting their representations. One of the most prominent approaches is structural probing (Hewitt and Manning, 2019), where a linear projection of word embeddings is performed in order to approximate the topology of dependency structures. In this work, we introduce a new type of structural probing, where the linear projection is decomposed into 1. isomorphic space rotation; 2. linear scaling that identifies and scales the most relevant dimensions. In addition to syntactic dependency, we evaluate our method on novel tasks (lexical hypernymy and position in a sentence). We jointly train the probes for multiple tasks and experimentally show that lexical and syntactic information is separated in the representations. Moreover, the orthogonal constraint makes the *Structural Probes* less vulnerable to memorization.

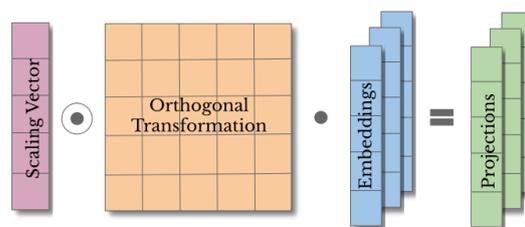
1 Introduction

Latent representations of neural networks encode specific linguistic features. Recently, a lot of focus was devoted to interpret these representations and analyze structures captured by the deep models. One of the most popular analysis methods is probing (Belinkov et al., 2017; Blevins et al., 2018; Linzen et al., 2016; Liu et al., 2019). The pre-trained model’s ¹ parameters are fixed, and its latent states or outputs are then fed into a simple neural network optimized to solve an auxiliary task, e.g., semantic, syntactic parsing, anaphora resolution, morphosyntactic tagging, etc. The amount of language information stored in the representations can be evaluated by measuring the specific language task’s performance.

¹Typically models for language modeling or machine translation are analyzed.



(a) *Structural Probe*



(b) *Orthogonal Structural Probe*

Figure 1: Comparison of the *Structural Probe* of Hewitt and Manning (2019) and the *Orthogonal Structural Probe* proposed by us.

Probing experiments usually involve classification tasks. Lately, Hewitt and Manning (2019) proposed *Structural Probes*, which use regression as an optimization objective. They train a linear projection layer to approximate: 1. dependency tree distances between words² by the Euclidean distance between transformed vectors; 2. the tree depth of a word by the norm of its vector.

In Figure 1, we visualize our *Orthogonal Structural Probe*. A linear transformation is replaced by an *Orthogonal Transformation* (rotation of the embedding space), and product-wise multiplication of rotated vectors by a *Scaling Vector* to get the final projections. Our motivation is to obtain an embedding space that is isomorphic with the original one, and the impact of each dimension can be evaluated

²Tree distance is the length of the tree path between two tokens

by analyzing *Scaling Vector*'s weights. We elaborate on mathematical properties and training details in Section 3.

In addition to dependency trees used by [Hewitt and Manning \(2019\)](#), we introduce new structural tasks related to lexical hypernymy and word's position in the sentence. We also employ a control task, in which we evaluate the memorization of randomly generated trees. *Orthogonal Structural Probes* let us optimize for multiple objectives jointly by keeping a shared *Orthogonal Transformation* matrix and changing task-specific *Scaling Vectors*.

We will answer the following questions:

1. Do our *Orthogonal Structural Probes* achieve comparable or better performance to the *Structural Probes* of [Hewitt and Manning \(2019\)](#)?
2. Can we find other phenomena such as lexical hypernymy and a word's absolute position in a sentence using *Orthogonal Structural Probe*? How vulnerable are the probes to memorizing random data?
3. Is it possible to effectively train *Orthogonal Structural Probes* jointly for multiple auxiliary objectives, i.e., depth and distance, or multiple types of structures mentioned in the previous question?
4. Can we identify particular dimensions of the embedding space that encode particular linguistic structures? Are there any superfluous dimensions?
5. If yes, what is the relationship between subspaces encoding distinct structures?

2 Related Work

Basic linguistic features can be easily extracted from the contextual representations ([Liu et al., 2019](#)). Probing was intensively used to investigate the representation of morphological information (mainly POS tags) in hidden states of machine translation systems and language models ([Belinkov et al., 2017](#); [Peters et al., 2018](#); [Tenney et al., 2019b](#)). Besides the work of [Hewitt and Manning \(2019\)](#), probing for dependency syntax was performed by [Tenney et al. \(2019a\)](#) and [Blevins et al. \(2018\)](#). They utilize a binary classifier to predict dependency edges. In work contemporary to ours,

[Ravichander et al. \(2020\)](#) employ a softmax classifier to show that BERT can be successfully probed for hypernymy.

There is an ongoing debate on which probe architectures offer a good insight into underlying representations. [Zhang and Bowman \(2018\)](#) showed that a POS tagger on top of a frozen randomly initialized LSTM model achieves unexpectedly high results. In the work of [Hewitt and Liang \(2019\)](#), the multilayer perceptron probes display similar accuracy for predicting POS tags as for randomly assigned tags. These symptoms underscore how crucial it is to carefully consider the probe's architecture to avoid reaching spurious conclusions. It is good practice to monitor additional aspects of the probe beyond performance on a linguistic task, such as selectivity ([Hewitt and Liang, 2019](#)), or complexity ([Pimentel et al., 2020](#)). The recent state of knowledge is summarized in surveys on probing ([Belinkov and Glass, 2019](#)) and interpretation of BERT's representations ([Rogers et al., 2020](#)).

Orthogonality has been applied broadly in the field of deep learning, especially to cope with exploding/vanishing gradient problem in recurrent neural networks ([Arjovsky et al., 2016](#); [Jing et al., 2017a](#); [Wisdom et al., 2016](#)). In this work, we use regularization to enforce the orthogonality of a dense layer. In literature, such an approach is called "soft constraint" ([Bansal et al., 2018](#); [Vorontsov et al., 2017](#)). Alternatively, "hard constraint" assumes parameterization of a network such that the transformation of latent states is orthogonal by definition ([Arjovsky et al., 2016](#); [Jing et al., 2017b](#)). There are a few examples of orthogonality applications in NLP: in RNN language model ([Dangovski et al., 2019](#)); in Performer ([Choromanski et al., 2020](#)), which is a more efficient counterpart of Transformer ([Vaswani et al., 2017](#)). Best to our knowledge, we are the first to use orthogonal transformation in probing.

3 Method

In this section, we first review the structural probing proposed by [Hewitt and Manning \(2019\)](#) and then introduce our *Orthogonal Structural Probe*.

3.1 Structural Probes

In the previous work, a linear transformation is optimized to transform the contextual word representations produced by a pre-trained neural model (e.g. BERT [Devlin et al. \(2019\)](#), ELMo [Peters et al.](#)

(2018)). The squared L2 norm of the differences between transformed word vectors approximate the tree distance between them:

$$d_B(h_i, h_j)^2 = (B(h_i - h_j))^T (B(h_i - h_j)), \quad (1)$$

where B is the *Linear Transformation* matrix and h_i, h_j are the vector representations of words at positions i and j .

The probe is optimized to approximate the distance between tokens in the dependency tree (d_T) by gradient descent objective:

$$\min_B \frac{1}{s^2} \sum_{i,j} |d_T(w_i, w_j) - d_B(h_i, h_j)|^2, \quad (2)$$

where s is the length of a sentence.

Moreover, the same work introduced depth probes, where vectors were linearly transformed so that the squared L2 length of the mapping approximate the token’s depth in a dependency tree:

$$\|h_i\|_B^2 = (Bh_i)^T (Bh_i) \quad (3)$$

Gradient descent objective is analogical:

$$\min_B \frac{1}{s} \sum_i \left| \|w_i\|_T - \|h_i\|_B^2 \right| \quad (4)$$

3.2 Orthogonal Structural Probes

We introduce orthogonality to structural probes. For that purpose, we perform the singular value decomposition of the matrix B

$$B = U \cdot D \cdot V^T, \quad (5)$$

where the matrices U and V are orthogonal, and D is diagonal. Notably, when we substitute B with $U \cdot D \cdot V^T$ in Eq. (1), the matrix U cancels out. It can be easily shown by rearranging the variables in the equation:³

$$\begin{aligned} d_B(h_i, h_j)^2 &= (DV^T(h_i - h_j))^T (DV^T(h_i - h_j)) \\ &= (DV^T(h_i - h_j))^T (D \odot V^T(h_i - h_j)) \end{aligned} \quad (6)$$

We can replace the diagonal matrix D with a vector \bar{d} and use element-wise product (we will call \bar{d} the *Scaling Vector*). Finally, we get the following equation for *Orthogonal Distance Probe*:

$$\begin{aligned} d_{\bar{d}V^T}(h_i, h_j)^2 &= (\bar{d} \odot V^T(h_i - h_j))^T (\bar{d} \odot V^T(h_i - h_j)) \end{aligned} \quad (7)$$

³A complete derivation can be found in the appendix.

The same reasoning can be applied to Eq. (3) to obtain *Orthogonal Depth Probe*:

$$\|h_i\|_{\bar{d}V^T}^2 = (\bar{d} \odot V^T h_i)^T (\bar{d} \odot V^T h_i) \quad (8)$$

We showed that *Orthogonal Structural Probe* is mathematically equivalent to *Standard Structural Probe*.

3.3 Multitask Training

Orthogonal Structural Probe can be easily adapted to multitask probing for a set of objectives \mathcal{O} . We use one shared *Orthogonal Transformation* and different *Scaling Vectors* for each task. In one batch, we compute a loss for a specific objective. For each batch (with objective $o \in \mathcal{O}$), a forward pass consists of multiplication by a shared orthogonal matrix V^T and product-wise multiplication by a designated vector \bar{d}_o . All the batches are shuffled together in a training epoch.

3.4 Orthogonality Regularization

We use *Double Soft Orthogonality Regularization* (DSO) proposed by Bansal et al. (2018) to coerce orthogonality of the matrix V during training:

$$\lambda_O DSO(V) = \lambda_O (\|V^T V - \mathbb{I}\|_F^2 + \|V V^T - \mathbb{I}\|_F^2) \quad (9)$$

$\|\cdot\|_F$ stands for the Frobenius norm of a matrix.

3.5 Sparsity Regularization

In further experiments, we investigate the effects of sparsity in *Scaling Vector*. For that purpose, we compute the L1 norm and add it to the training loss.

$$\lambda_S \|\bar{d}\|_1 \quad (10)$$

3.6 Training Objective

Altogether, the loss equation in *Orthogonal Distance Probe* for objective $o \in \mathcal{O}$ is the following:

$$\begin{aligned} L_{o, dist.} &= \frac{1}{s^2} \sum_{i,j} |d_T(w_i, w_j) - d_{\bar{d}_o V^T}(h_i, h_j)|^2 + \\ &\quad + \lambda_O DSO(V) + \lambda_S \|\bar{d}_o\|_1 \end{aligned} \quad (11)$$

And in *Orthogonal Depth Probe*:

$$\begin{aligned} L_{o, depth} &= \frac{1}{s} \sum_i \left| \|w_i\|_T - \|h_i\|_{\bar{d}_o V^T}^2 \right| + \\ &\quad + \lambda_O DSO(V) + \lambda_S \|\bar{d}_o\|_1 \end{aligned} \quad (12)$$

The loss is normalized by the number of predictions in a sentence and averaged across a batch.

4 Experiments

We train probes on top of each of 24 layers of English BERT large cased model (Devlin et al., 2019) implemented by HuggingFace (Wolf et al., 2020). We optimize for the approximation of depth and distance in four types of structures: syntactic dependency, lexical hypernymy, absolute position in a sentence, and randomly generated trees. In the following subsection, we expand upon these structures.

4.1 Data and Objectives

In our experiments, we use training, evaluation, and test sentences from Universal Dependencies English Web Treebank (Silveira et al., 2014). Depending on the objective, we reveal only partial relevant annotation from the dataset.

Dependency Syntax We probe for syntactic structure in Universal Dependencies parse trees (Nivre et al., 2020). Dependency trees are annotated in English Web Treebank. We focus on distances between words in dependency trees and their depth, i.e., distance from the syntactic root.

Lexical Hypernymy We introduce probing for lexical information. We optimize probes to approximate the distance between pairs of words in the hypernymy tree and the depth for each word. For that purpose, we use the tree from WordNet (Miller, 1995). We consider lexical distances between pairs of nouns and pairs of verbs in sentences and lexical depth for each noun and verb. We provide gold POS information and look up synset by a lemmatized form of a word to avoid ambiguity.

Position in a Sentence Probing for the sentence index of a word and positional difference between pairs of words.

Random Structures We probe for randomly generated trees. When we jointly optimize for depth and distance, we keep the same randomly generated tree. This control task allows us to determine the extent to which our probes memorize the structures and thus over-fit to the training data.

4.2 Training

We use batches of size 12 and an initial training rate of 0.02. We use learning rate decay and early-stopping mechanism: if validation loss does not achieve a new minimum after an epoch, the learning rate is divided by 10. After three consecutive

learning rate updates not resulting in a new minimum, the training is stopped.

Orthogonality Regularization In our experiments, we took λ_O equal to 0.05.⁴ The regularization converged early during the gradient optimization. Hence we can assume that matrix V is orthogonal.

Sparsity Regularization By default $\lambda_S = 0$. Only in the experiments described in Section 5.1, we use sparsity regularization by setting λ_S to a positive value (0.005, 0.05, or 0.1) when DSO drops below 1.5 during the training. This mechanism prevents weakening orthogonality constraint in early epochs.

Additional details of the training are described in the appendix. The code is available at GitHub: <https://github.com/Tom556/OrthogonalTransformerProbing>.

4.3 Evaluation

We assess Spearman’s rank correlation between gold and predicted values. We report the average correlations for the sentences with lengths from 5 to 50 in the same way as Hewitt and Manning (2019).

Our *Orthogonal Structural Probes* are trained jointly for multiple objectives (Section 3.3). We evaluate the effect of multitasking testing different configurations: **A)** separate probing for each objective; **B)** joint probing for distance and depth in the same structure type; **C)** joint probing for distance in all structures; **D)** joint probing for depths in all structures; **E)** probing for all objectives together. We compare the results with two baselines: **I)** optimizing only *Scaling Vector*; **II)** *Structural Probes*.

4.4 Dimensionality of Scaling Vector

We hypothesize that the orthogonality regularization allows us to find embedding subspace capable of representing a particular linguistic structure. In Section 5.1, we examine the performance of lower-rank projections and ask whether further restrictions of dimensionality affect the results. In Section 5.2 we analyze interactions between subspaces related to a particular objective in a joint probing setting.

⁴We experimentally checked that ten times smaller and ten times larger values of λ_O do not affect orthogonality of matrix V and lead to the same results.

	I	II	A	B	C / D	E
	Scaling Vector only	Structural Probe	Orthogonal Structural Probe	multitask orthogonal probing		
				distance + depth	all distances or all depths	all tasks
DEP Depth Layer	.459 \pm .001 17	.856 \pm .001 18	<u>.858</u> \pm .001 17	.855 \pm .001 16	.850 \pm .002 16	.852 \pm .001 16
DEP Dist. Layer	.513 \pm .001 18	<u>.843</u> \pm .001 17	.842 \pm .001 17	.838 \pm .001 17	.833 \pm .001 17	.832 \pm .002 16
LEX Depth Layer	.572 \pm .001 13	<u>.892</u> \pm .002 8	.882 \pm .002 8	.869 \pm .005 8	.885 \pm .004 6	.873 \pm .005 9
LEX Dist. Layer	.560 \pm .001 13	<u>.816</u> \pm .008 6	.803 \pm .005 6	.789 \pm .004 7	.792 \pm .010 6	.792 \pm .005 6
POS Depth Layer	.232 \pm .013 5	<u>.989</u> \pm .001 1	.983 \pm .001 6	.986 \pm .001 1	.976 \pm .004 2	.982 \pm 0.001 3
POS Dist. Layer	.441 \pm 0.001 1	<u>.980</u> \pm .001 4	.979 \pm .001 4	.977 \pm .001 4	.978 \pm .001 5	.976 \pm 0.001 4
RAND Depth Layer	.008 \pm .002 6	.206 \pm .010 17	.136 \pm .007 18	.129 \pm .010 18	.163 \pm .023 18	<u>.107</u> \pm .019 19
RAND Dist. Layer	.149 \pm .001 17	.242 \pm .005 19	.220 \pm .006 18	<u>.206</u> \pm .004 17	.209 \pm .005 19	.208 \pm .007 15
AVG. DEP, LEX, POS ABOVE - AVG. RAND	.463 .385	.896 .673	.891 .713	.886 .718	.886 .699	.883 .726

Table 1: The highest Spearman’s correlations (across layers) between predicted values and gold annotations on a held out test set (for random structures computed on a train set). Each column represents another variant of training. Standard deviation was calculated for six runs. Each row’s optimal result is underlined (except baseline I); results within 95% confidence interval based on Student’s t-test (Student, 1908) are marked in bold.

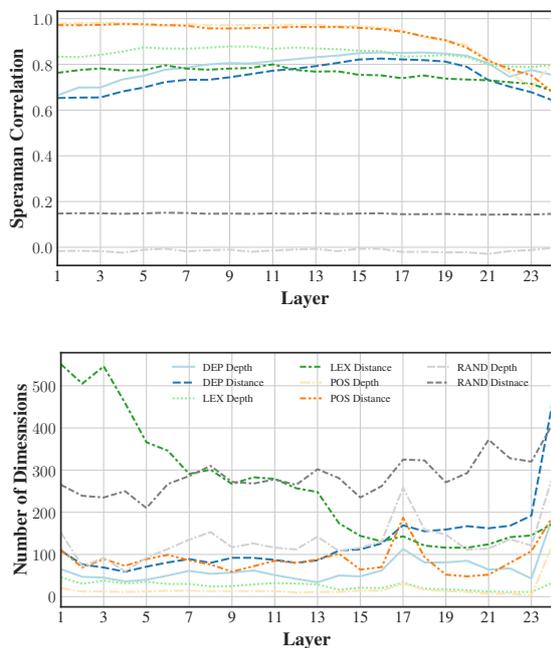


Figure 2: Spearman correlations and number of non-zero *Scaling Vector*’s dimensions across layers for joint training.

5 Results

We compare Spearman’s correlations between predicted values and gold tree depths and distances in Table 1. The correlations obtained from *Orthogonal Structural Probes* are high for linguistic structures: from 0.803 for lexical distance to 0.882 for lexical depth. Predicted positional depths and distances nearly match gold values.

Correlation on training data for random structures is very weak, hinting that the probes do not memorize structures during training but extract them from the model’s representations. The correlation for distances is higher than for depth. We hypothesize it is because the probes learn some basic tree properties.⁵

The results obtained by *Orthogonal Structural Probes* are close to those of *Structural Probes*. For dependency distance, the difference is not statistically significant. Notably, correlations on training set for randomly generated trees decreased. It suggests that *Orthogonal Structural Probes* are less vulnerable to memorization. In multitask probing,

⁵For instance, when the distances between nodes X and Y, and Y and Z are both 1, then the distance between X and Z needs to be 2

	Subspace		Share of Dropped Dimensions			Sparsity Regularization					
	Dims	Corr	25%	33%	50%	$\lambda_S = 0.005$		$\lambda_S = 0.05$		$\lambda_S = 0.1$	
						Dims	Corr	Dims	Corr	Dims	Corr
DEP Depth	137	.858	.783	.758	.700	26	.856	2	.832	1	.822
DEP Dist.	189	.842	.800	.781	.741	76	.835	21	.784	14	.746
LEX Depth	19	.884	.841	.822	.784	19	.875	11	.852	10	.836
LEX Dist.	263	.805	.768	.755	.722	92	.792	60	.756	52	.737
POS Depth	20	.983	.760	.686	.526	11	.982	6	.981	3	.981
POS Dist.	98	.979	.890	.859	.627	38	.978	14	.975	11	.970
RAND Depth	259	.128	.108	.101	.091	6	.037	1	.011	1	.010
RAND Dist.	399	.222	.215	.213	.208	116	.208	20	.163	13	.155

Table 2: The highest Spearman’s correlations (across layers) between predicted values and gold annotations on a held-out test set (for random structures computed on a train set). In columns 2-3, results, when only selected dimensions are used. In columns 4-6, a portion of the selected dimensions is masked. In columns 7-12, sparsity regularization with different λ_S is applied. Probing for one objective.

correlation evenly decreases across all tasks. While selectivity (the difference between average correlation for dependency, lexical, and positional objectives and random objectives) increases from 0.673 to 0.726. Optimizing only a *Scaling Vector* gives distinctly lower correlations. These results emphasize the necessity of changing the coordinate system to amplify the dimensions encoding linguistic information.

In Fig. 2 (upper), we observe that the performance varies throughout the layers, confirming previous observations by Hewitt and Manning (2019) and Tenney et al. (2019a). The mid-upper layers tend to be more syntactic, and the mid-lower ones are more lexical. Predicting word position is more accurate in the lower layers, dropping significantly toward the last layers. It is due to the fact that in BERT, positional embeddings are added before the first layer. Random structure probes maintain steady results across all the layers.

5.1 Dimensionality

We observe that orthogonality constraint is quite effective in restricting the probe’s rank. In most of our experiments, the majority of *Scaling Vector* parameters converged to zero. It allows selecting subspaces encoding particular linguistic features. We want to answer whether such subspace has enough capacity for each probing task. For that purpose, we zero out the dimensions with corresponding *Scaling Vector* weights closer to zero

than $\epsilon = 10^{-4}$.⁶ Their elimination does not affect the results; correlations in Table 2 and Table 1 column A are practically equal. The dimensionality reduction is the strongest for lexical and positional depth probes, where subspaces with the rank of 19 and 20 respectively encode the structures as well as the whole embedding space with 1024 dimensions (Fig. 2, lower). The number of selected dimensions is the highest in probing for random structures. This is because a large capacity is required for memorization.

Another question we pose is whether it would be adequate to shrink the subspace even further. For each objective, we choose and drop a random portion of parameters to examine how it would affect the predictions. We conduct a procedure similar to cross-validation, i.e., we repeatedly drop disjoint and exhaustive sets of dimensions and average results for each set at the end.⁷ Table 2 shows that dimension dropping had the largest impact on positional probes: -0.458 for depth; the decrease is low for lexical distance – only -0.083 . It suggests that the information necessary for the latter objective is more dispersed than for the former one.

Sparsity Regularization We use sparsity regularization of *Scaling Vector* to examine whether dimensionality can be reduced more intelligently. The strength of regularization is regulated by value

⁶In the appendix, we show that dimension selection is not sensitive to the selection of low $10^{-30} < \epsilon < 10^{-3}$.

⁷When we drop 25% of dimensions, we randomly choose four sets. Each dimension is exactly in one set.

		DEP		LEX		POS		RAND	
		Depth	Dist.	Depth	Dist.	Depth	Dist.	Depth	Dist.
DEP	Depth	62	48	0	0	10	19	23	21
	Dist.		126	0	0	9	23	25	30
LEX	Depth			20	18	0	4	1	5
	Dist.				131	0	7	5	19
POS	Depth					14	10	13	10
	Dist.						70	33	50
RAND	Depth							131	95
	Dist.								262

Table 3: The number of shared dimensions selected by *Scaling Vector* after the joint training of probe on top of the 16th layer.

of $\lambda_s \in \{0.005, 0.05, 0.1\}$. We observe that for some objectives (dependency depth, positional depth, and positional distance), the relevant information is captured in a small number of dimensions. Remarkably, only one dimension of embedding space can achieve 0.822 correlation with dependency depths. We conjecture that if it is possible to achieve a high correlation with sparse subspaces, information on the phenomenon is focal in the model (concentrated in few dimensions). For the objectives with focal information, results decrease sharply when random dimensions are dropped because the probability of dropping important coordinates is high. On the other end of the spectrum, we can identify the objective for which information is spread – lexical distance. The dropping of random dimensions only moderately decreases correlation, as there are no especially essential coordinates. Probing with sparsity regularization produces subspaces of relatively large size.

Sparsity regularization also positively affects control objectives, decreasing correlations with distances and depths of randomly generated structures, indicating that regularized probes are less prone to memorization.

Notably, [Torroba Hennigen et al. \(2020\)](#) proposed a method for selecting embeddings’ dimensions relevant to particular linguistic phenomena. In our setting, thanks to the *Orthogonal Transformation*, we are not constrained to analyzing the dimensions of just one coordinate system.

5.2 Separation of Information

Another outcome of joint training was the ability to examine relationships between subspaces for each of the objectives. Figure 3 shows histograms of the dimensions selected in lexical and dependency probes. Each bin of the histogram corresponds to 10 coordinates. The height of a bar (in one color) represents how many were selected for a specific task. The dimensions on the x-axis are ordered by the weighted absolute values of *Scaling Vectors*.⁸

We found that in layers 6 and 16 (they achieve the highest correlation in lexical and dependency, respectively), the histograms are disjoint, indicating that the layers’ representations of dependency syntax and lexical hypernymy are orthogonal to each other in the embedding space. The orthogonality is less visible in the first layer and disappears almost entirely in the top one. In most layers, depth subspace is included in distance subspace for the same structural type. This behavior was expected as distance probing is more complex and therefore requires more capacity.

In Fig. 4 we present histograms for additional tasks at the model’s 16th layer. The positional subspace has a sizable intersection with the syntactic one, yet only a few common dimensions with the lexical subspace. The connection can be attributed to the fact that dependency edges can often be inferred from words’ relative positions. Probing for random structures is interlinked with other objectives. The sizes of shared subspaces for each pair can be found in Table 3. Histograms and tables for other sets of tasks are presented in the appendix.

6 Discussion

The introduction of an orthogonal constraint is a core element of our analysis. The constraint assures that no dimension is enhanced or diminished in the transformation and allows interpreting the magnitude of values in the *Scaling Vector* as the relevance of each dimension for the objectives.

In an *Orthogonal Structural Probe*, the sufficient rank of a transformation is learned during the optimization. The rank regularization is a prerequisite to disentangle the information encoded by the probe (Section 5.2). The natural question

⁸We weight the values before sorting to keep together non-zero dimensions of each *Scaling Vector*, i.e., dependency depth values are multiplied by 1000, dependency distance 100, lexical depth by 10. The weighting is performed only for visualization; the separation of linguistic information can be observed independently in Table 3.

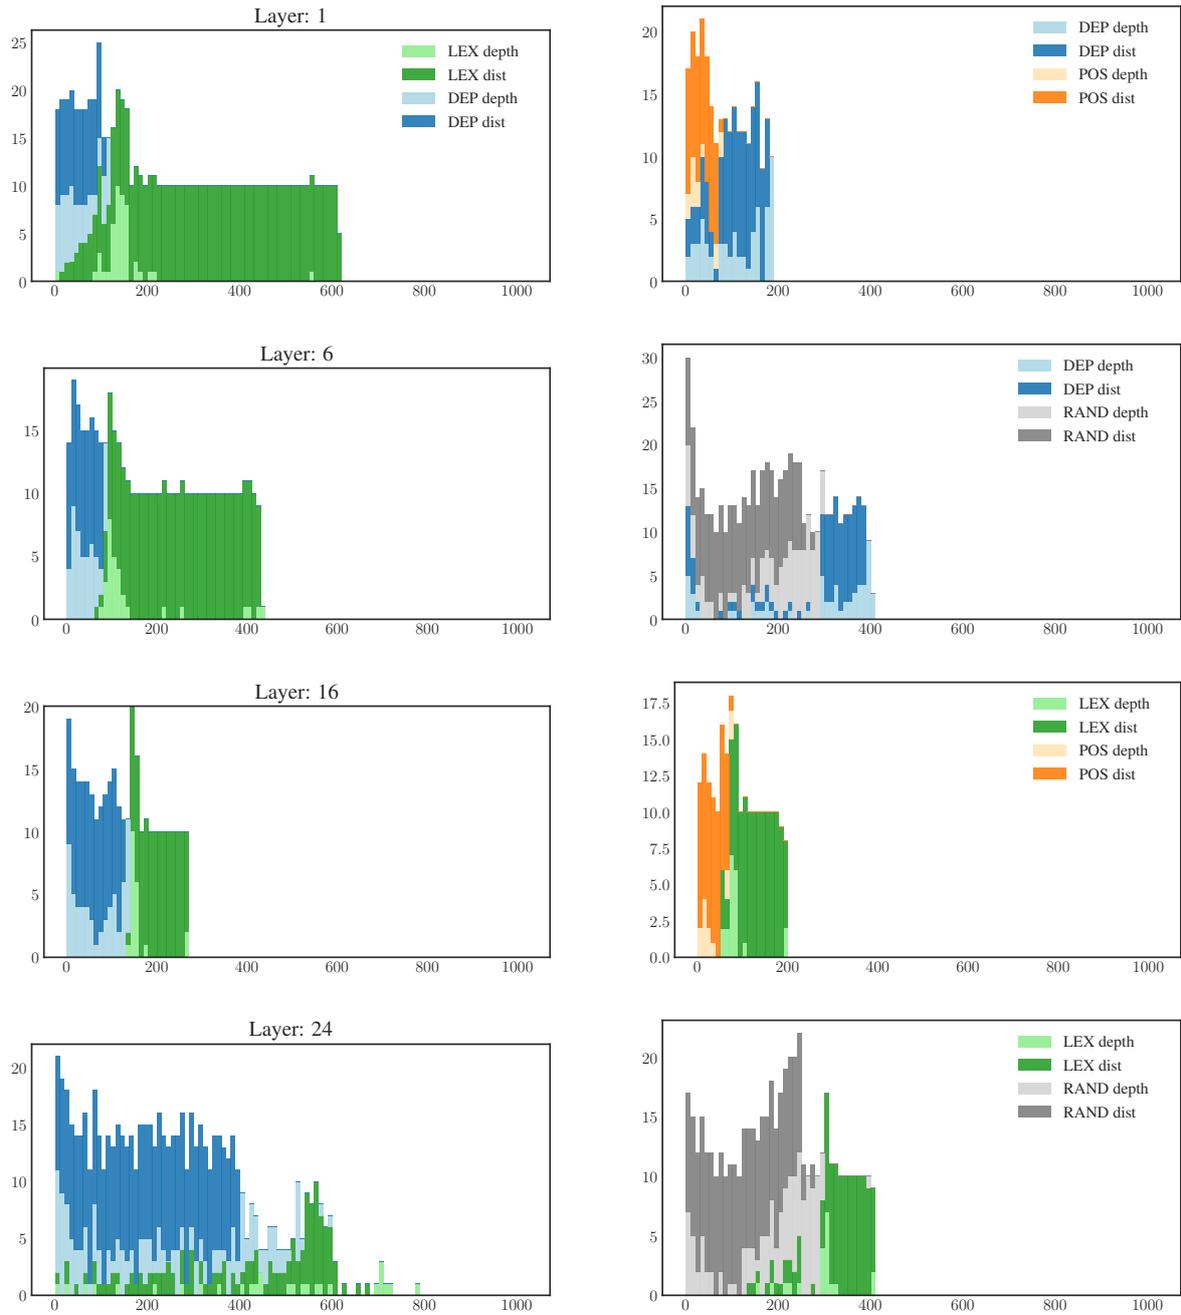


Figure 3: Histograms of dimensions selected by dependency and lexical *Scaling Vector* after joint training . Best in color.

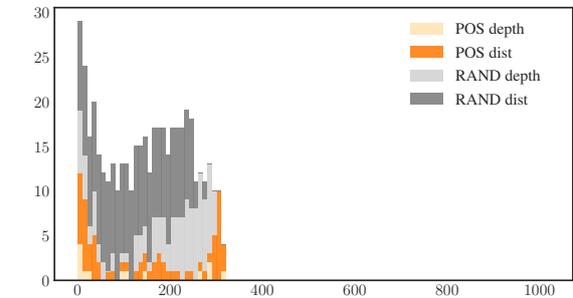


Figure 4: Histograms of dimensions selected by *Scaling Vector* after the joint training of probe on top of the 16th layer. Best in color.

is whether such analysis can be performed by reducing the rank of *Structural Probe* with another regularizer and decomposing linear transformation after the optimization. We argue that it is not possible both in joint and separate probing:

- In joint probing for multiple tasks: one *Scaling Vector* is shared for all the tasks. It is not possible to attribute the dimensions to a specific task.
- In separate probing for each task: the decomposition leads to different orthogonal matrices. Hence, the dimensions of distinct *Scaling Vectors* do not correspond to each other.

6.1 Limitations

We focus on syntax annotated in Universal Dependencies and lexical hypernymy encoded in WordNet. We do not claim that there is no correlation between syntactic and lexical information in BERT, just that the topologies of those two structures are encoded separately. It is entirely possible that we could find dimensions overlap when probing for syntax and lexicon in differently annotated datasets.

Conversely to *Structural Probes*, our reformulation of the loss (in Eq. (12) and Eq. (11)) is not convex. We thank one of the anonymous ACL reviewers for pointing it out. Nevertheless, we show that despite non-convexity, our *Orthogonal Structural Probes* achieve similar results to *Structural Probes* and are more selective.

7 Conclusions

We have expanded structural probing to new types of auxiliary tasks and introduced a new setting, *Orthogonal Structural Probe*, in which probes can be optimized jointly. We found out that:

1. Results of *Orthogonal Structural Probes* are on par with *Standard Structural Probes* on linguistic tasks. *Orthogonal Structural Probes* are less vulnerable to memorization.
2. In addition to syntactic dependencies *Orthogonal Structural Probes* can be efficiently trained to approximate dependency and depth in WordNet hypernymy trees and positional order.
3. *Orthogonal Structural Probes* can be trained jointly for multiple objectives. In most cases,

the performance moderately drops, and selectivity increases. The number of parameters decreases in comparison to training many separate probes.

4. Usually, information necessary for each objective is stored in a subspace of relatively low rank (19 - 263). We can further reduce dimensionality by applying sparsity regularization. For a few objectives (e.g., positional depth, dependency depth), the information is hugely focal, and the performance can fall markedly when just 25% randomly selected dimensions are dropped.
5. We have found that in most of BERT's layers, the subspace encoding linguistic hypernymy is separated from the subspace encoding dependency syntax and subspace encoding word's position.

7.1 Further work

Our method can be adjusted for multitask and multilingual settings. Following the observation that the orthogonal transformation can map distributions of embeddings in typologically close languages (Mikolov et al., 2013; Vulić et al., 2020). We think that joint training for many languages may be possible by keeping the same *Scaling Vector* and adding a separate *Orthogonal Transformation* per language, fulfilling the role of orthogonal mappings. Another leg of research would be analyzing probes for other linguistic structures, for instance, derivation trees.

Acknowledgments

We thank Ondřej Dušek, Greg Durrett, and anonymous reviewers of ACL for valuable comments on previous versions of this paper. This work has been supported by grant 18-02196S of the Czech Science Foundation and by grant 338521 of the Charles University Grant Agency. We have been using language resources and tools developed, stored and distributed by the LINDAT/CLARIAH-CZ project of the Ministry of Education, Youth and Sports of the Czech Republic (project LM2018101).

References

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp,

- Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. [TensorFlow: Large-scale machine learning on heterogeneous systems](#). Software available from tensorflow.org.
- Martin Arjovsky, Amar Shah, and Yoshua Bengio. 2016. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pages 1120–1128.
- Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. 2018. [Can We Gain More from Orthogonality Regularizations in Training Deep Networks?](#) In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 4261–4271. Curran Associates, Inc.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey](#). *Transactions of the Association for Computational Linguistics (TACL)*, 7:49–72.
- Terra Blevins, Omer Levy, and Luke Zettlemoyer. 2018. [Deep RNNs encode soft hierarchical syntax](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 14–19, Melbourne, Australia. Association for Computational Linguistics.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarnos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. 2020. [Rethinking attention with performers](#).
- Rumen Dangovski, Li Jing, Preslav Nakov, Mićo Tat-alović, and Marin Soljačić. 2019. Rotational unit of memory: a novel representation unit for RNNs with scalable applications. *Transaction of the Association of Computational Linguistics*, 7:121–138.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- 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.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. In *NAACL-HLT*.
- Li Jing, Caglar Gulcehre, John Peurifoy, Yichen Shen, Max Tegmark, Marin Soljačić, and Y. Bengio. 2017a. [Gated orthogonal recurrent units: On learning to forget](#). *Neural Computation*, 31.
- Li Jing, Yichen Shen, Tena Dubcek, John Peurifoy, Scott Skirlo, Yann LeCun, Max Tegmark, and Marin Soljačić. 2017b. [Tunable efficient unitary neural networks \(EUNN\) and their application to RNNs](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1733–1741, International Convention Centre, Sydney, Australia. PMLR.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Artur Kulmizev, Vinit Ravishankar, Mostafa Abdou, and Joakim Nivre. 2020. [Do neural language models show preferences for syntactic formalisms?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4077–4091, Online. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. Linguistic knowledge and transferability of contextual representations. In *NAACL-HLT*.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation.
- George A. Miller. 1995. [Wordnet: A lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.

- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, New Orleans, Louisiana. Association for Computational Linguistics.
- 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.
- Abhilasha Ravichander, Eduard Hovy, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. [On the systematicity of probing contextualized word representations: The case of hypernymy in BERT](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 88–102, Barcelona, Spain (Online). Association for Computational Linguistics.
- 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*, 8:842–866.
- Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.
- Student. 1908. The probable error of a mean. *Biometrika*, pages 1–25.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Lucas Torroba Hennigen, Adina Williams, and Ryan Cotterell. 2020. [Intrinsic probing through dimension selection](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 197–216, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Eugene Vorontsov, Chiheb Trabelsi, Samuel Kadoury, and Chris Pal. 2017. [On orthogonality and learning recurrent networks with long term dependencies](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3570–3578, International Convention Centre, Sydney, Australia. PMLR.
- Ivan Vulić, Sebastian Ruder, and Anders Søgaard. 2020. [Are all good word vector spaces isomorphic?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3178–3192, Online. Association for Computational Linguistics.
- Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. 2016. [Full-capacity unitary recurrent neural networks](#). In *Advances in Neural Information Processing Systems*, volume 29, pages 4880–4888. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Kelly Zhang and Samuel Bowman. 2018. [Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361, Brussels, Belgium. Association for Computational Linguistics.

A Technical Details

The *Orthogonal Structural Probe* is trained to minimize L1 loss between predicted and gold distances and depths. The loss is normalized by the number of predictions in a sentence and averaged across a batch of size 12. Optimization is conducted with Adam (Kingma and Ba, 2014) with initial learning rate 0.02 and meta parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. We use learning rate decay and early-stopping mechanism: if validation loss does not achieve a new minimum after an epoch, learning rate is divided by 10. After three consecutive learning rate updates not resulting in a new minimum, the training is stopped.

To alleviate sharp jumps in training loss that we observed mainly in training of *Depth Probes*, we clip each gradient’s norm at $c = 1.5$.

We implemented the network in TensorFlow 2 (Abadi et al., 2015). The code is available at GitHub: <https://github.com/Tom556/OrthogonalTransformerProbing>.

A.1 Orthogonal Regularization

In order to coerce orthogonality of matrix V we add DSO to the loss. Bansal et al. (2018) showed that for convolutional neural network applied to image processing, a simpler regularization – SO is more powerful.

$$\lambda_{SO}SO(V) = \lambda_{SO}\|V^TV - \mathbb{I}\|_F^2 \quad (13)$$

In our experiments, DSO led to faster convergence. Fig. 5 shows values of orthogonality penalty during the training. Taking into account the properties of the Frobenius norm, we observe that V matrix is close to orthogonal already after initial epochs.

A.2 Sparsity Regularization

Fig. 6 presents values of sparsity penalty during the training. The regularization is applied only after the orthogonality penalty drops below 1.5.

A.3 Number of Parameters

The number of *Orthogonal Structural Probe*’s parameters is given by equation:

$$NParams_{Ortho} = D_{emb}^2 + D_{emb} \cdot N_{obj}, \quad (14)$$

where D_{emb} is dimensionality of the embeddings and N_{obj} is a number of jointly probed objectives. Therefore, our biggest probes on top of

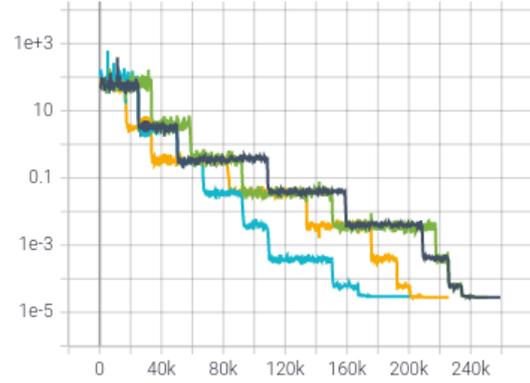


Figure 5: Values of orthogonality penalty during joint training of *Orthogonal Structural Probe* on top of layers: 3 (green), 7 (yellow), 16 (gray), 24 (blue). Optimization steps on the x-axis.

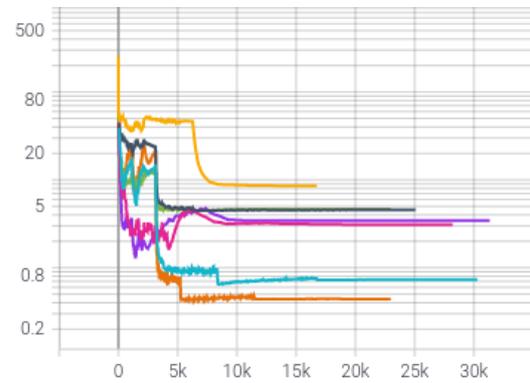


Figure 6: Values of sparsity penalty during separate training of *Orthogonal Structural Probes* with $\lambda = 0.05$. Objectives from the highest to the lowest value: lexical distance (yellow), positional distance (green), dependency distance (gray), positional depth (violet), lexical depth (magenta), dependency depth (blue), random depth (orange). Optimization steps on the x-axis.

BERT Large for all eight objectives have $1024^2 + 1024 \cdot 8 = 1,056,768$ parameters. It is more than in *Structural Probes* of Hewitt and Manning (2019). Nevertheless, our probes have less degrees of freedom, because we use *Orthogonal Transformation* instead of *Linear Transformation*.

$$DoF_{Ortho} = \frac{D_{emb} \cdot (D_{emb} - 1)}{2} + D_{emb} \cdot N_{obj} \quad (15)$$

In the case of joint training for all objectives, the number of degrees of freedom equals to 523,766.

A.4 Computation Time

We have trained *Orthogonal Structural Probes* on GPU a core *GeForce GTX 1080 Ti*. Approximate run times of specific configurations:

- separate probing for depth ~ 3 minutes
- separate probing for distance ~ 5 minutes
- joint probing for distance and depth in the same structure type ~ 7 minutes
- joint probing for depths in all structures ~ 13 minutes
- joint probing for distance in all structures ~ 18 minutes
- probing for all objectives together ~ 35 minutes

B Derivation of Orthogonal Structural Probe Equation

Eq. (6) with intermediate steps:

$$\begin{aligned}
 d_B(h_i, h_j)^2 &= (UDV^T(h_i - h_j))^T(UDV^T(h_i - h_j)) \\
 &= (h_i - h_j)^TVD^T U^TUDV^T(h_i - h_j) \quad (16) \\
 &= (h_i - h_j)^TVD^TDV^T(h_i - h_j) \\
 &= (DV^T(h_i - h_j))^T(DV^T(h_i - h_j))
 \end{aligned}$$

C Dataset Description

Universal Dependencies English Web Treebank (Silveira et al., 2014) is available at https://github.com/UniversalDependencies/UD_English-EWT. It consist of: 12,543 test, 2,002 dev, and 2,077 test sentences.

D Application in Dependency Parsing

We have computed the UAS of dependency trees predicted based on dependency probes. We employ the algorithm for extraction of directed dependency trees proposed by Kulmizev et al. (2020). Our innovation to the method is that we optimize distance and depth probes jointly during one optimization.

In line with the previous studies, we show that *Orthogonal Structural Probes* can be employed for parsing. Table 4 presents Unlabeled Attachment Scores achieved by different multi-task configurations. Joint probing for dependency distance and depth allows us to extract a directed dependency tree in just one optimization. Best to our knowledge, it has not been tried before. Analogically to Spearman’s correlation, UAS drops when more objectives are used in optimization. However, even joint probing for all eight objectives is capable of producing trees with 75.66% UAS.

Training config.	Layer	UUAS	UAS
Structural Probe	15	82.29	–
Orthogonal Probe	15	82.47	–
multitask orthogonal probing			
distance + depth	16	80.86	77.51
all distances	15	80.72	–
all tasks	16	79.03	75.66

Table 4: (Undirected) Unlabeled Attachment Score of trees extracted from dependency probes.

E Scaling Vector Properties

In this appendix, we elaborate on the properties of *Scaling Vectors* parameters in the multi-task probing.

E.1 Parameters Distribution

The distribution of values in *Scaling Vector* (Fig. 7) shows that the majority of parameters converge to zero. They are within 10^{-40} to 10^{-30} margin after training. Therefore, the significant dimensions are clearly identifiable.

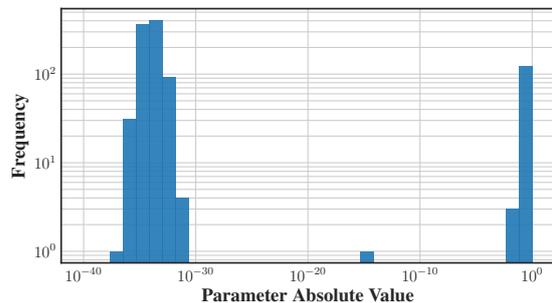
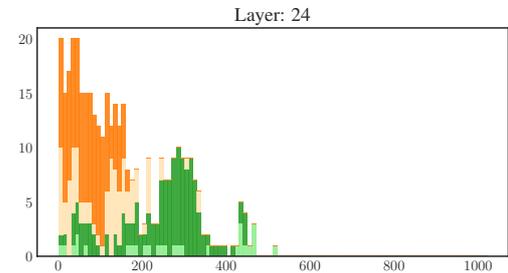
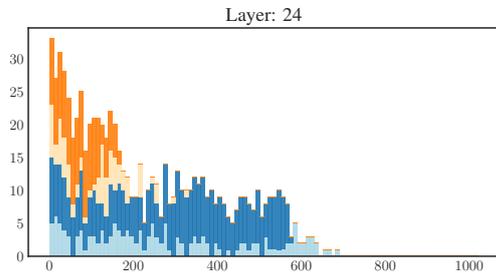
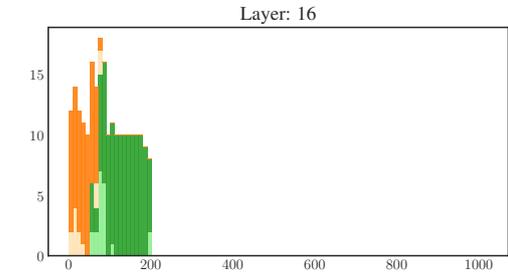
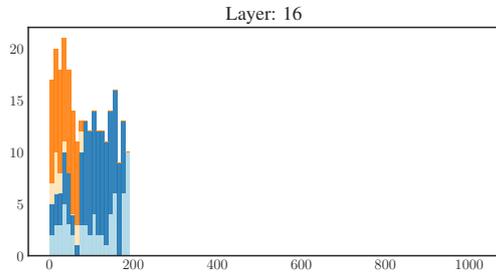
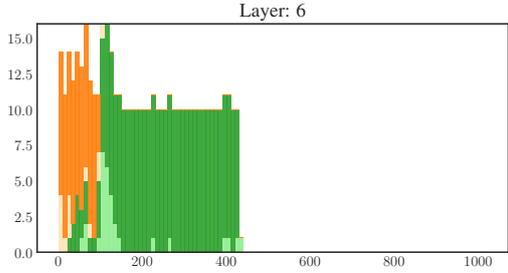
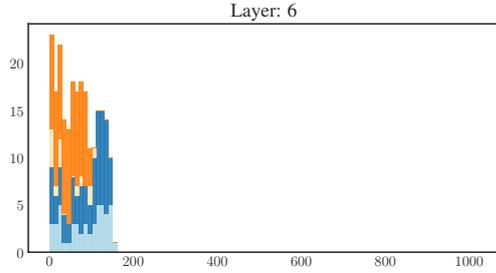
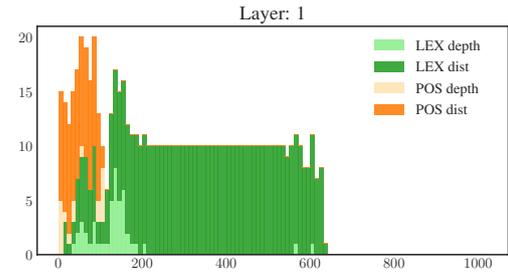
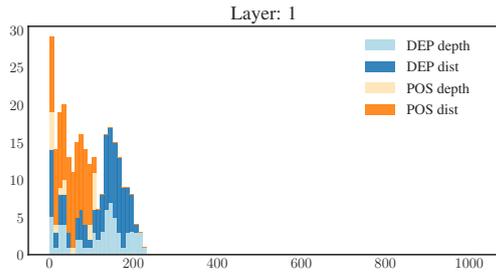


Figure 7: Logarithmic histogram of *Scaling Vector* parameters for dependency distance. Joint probing of 16th layer’s representations.

E.2 Separation of Information (Continued)

On the following pages, we present dimension overlap histograms and tables, as in Section 5.2, for the remaining pairs of objectives.

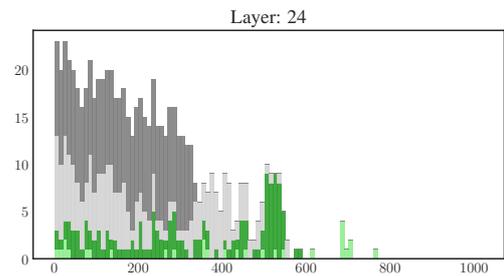
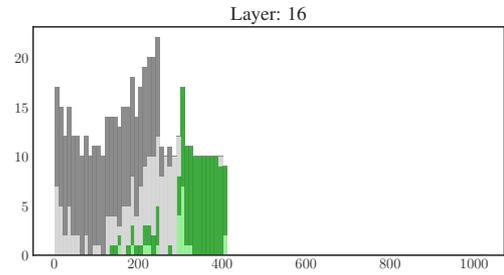
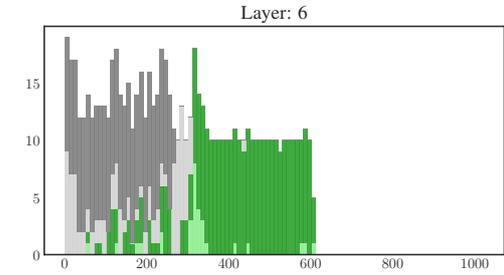
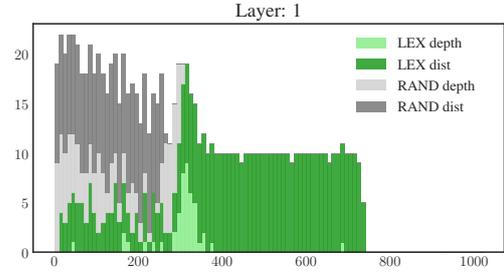
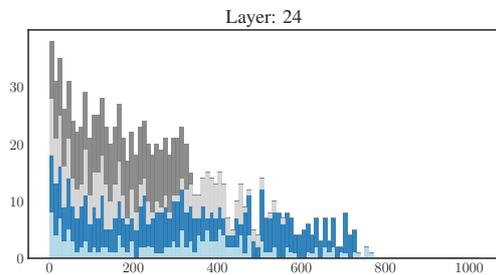
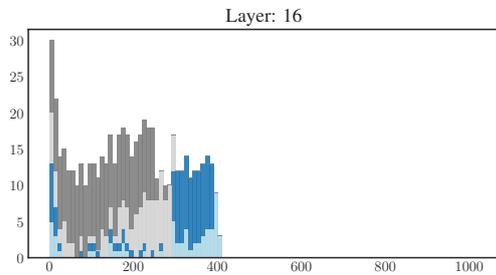
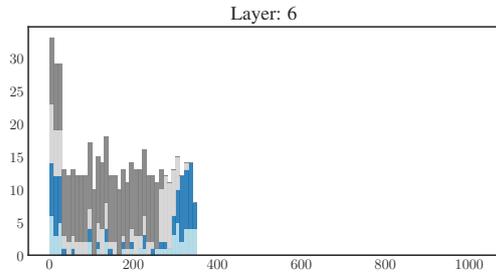
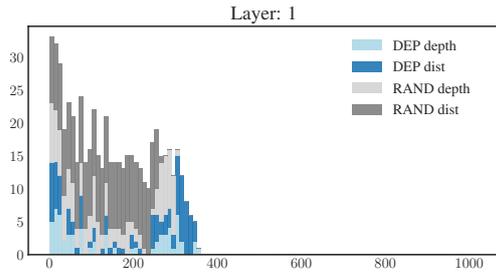


		DEP		LEX		POS		RAND	
		Depth	Dist.	Depth	Dist.	Depth	Dist.	Depth	Dist.
DEP	Depth	65	54	1	18	11	24	48	44
	Dist.		109	6	43	11	39	45	64
LEX	Depth			46	45	2	13	7	8
	Dist.				551	2	42	46	103
POS	Depth					20	11	20	14
	Dist.						111	47	71
RAND	Depth							152	112
	Dist.								265

		DEP		LEX		POS		RAND	
		Depth	Dist.	Depth	Dist.	Depth	Dist.	Depth	Dist.
DEP	Depth	50	43	0	1	11	26	30	26
	Dist.		81	1	2	11	38	35	39
LEX	Depth			30	28	0	4	1	6
	Dist.				346	0	19	14	45
POS	Depth					14	11	13	11
	Dist.						99	41	71
RAND	Depth							113	70
	Dist.								267

Table 5: Number of shared dimensions selected by *Scaling Vector* after the joint training of probe on top of the 1st layer.

Table 6: Number of shared dimensions selected by *Scaling Vector* after the joint training of probe on top of the 6th layer.



		DEP		LEX		POS		RAND	
		Depth	Dist.	Depth	Dist.	Depth	Dist.	Depth	Dist.
DEP	Depth	189	144	17	39	70	66	146	123
	Dist.		463	16	82	81	141	186	275
LEX	Depth			33	22	9	10	18	16
	Dist.				173	30	48	64	98
POS	Depth					124	70	107	97
	Dist.						190	136	177
RAND	Depth							287	198
	Dist.								410

Table 7: Number of shared dimensions selected by *Scaling Vector* after the joint training of probe on top of the 24th layer.