

A Non-Linear Structural Probe

Jennifer C. White^δ Tiago Pimentel^δ Naomi Saphra^ϑ Ryan Cotterell^{δ,ϒ}
^δUniversity of Cambridge, ^ϑUniversity of Edinburgh, ^ϒETH Zürich

jw2088@cam.ac.uk, tp472@cam.ac.uk
n.saphra@ed.ac.uk, ryan.cotterell@inf.ethz.ch

Abstract

Probes are models devised to investigate the encoding of knowledge—e.g. syntactic structure—in contextual representations. Probes are often designed for simplicity, which has led to restrictions on probe design that may not allow for the full exploitation of the structure of encoded information; one such restriction is linearity. We examine the case of a structural probe (Hewitt and Manning, 2019), which aims to investigate the encoding of syntactic structure in contextual representations through learning only linear transformations. By observing that the structural probe learns a metric, we are able to kernelize it and develop a novel non-linear variant with an identical number of parameters. We test on 6 languages and find that the radial-basis function (RBF) kernel, in conjunction with regularization, achieves a statistically significant improvement over the baseline in all languages—implying that at least part of the syntactic knowledge is encoded non-linearly. We conclude by discussing how the RBF kernel resembles BERT’s self-attention layers and speculate that this resemblance leads to the RBF-based probe’s stronger performance.

1 Introduction

Probing has been widely used in an effort to better understand what linguistic knowledge may be encoded in contextual word representations such as BERT (Devlin et al., 2019) and ELMo (Peters et al., 2018). These probes tend to be designed with simplicity in mind and with the intent of revealing what linguistic structure is encoded in an embedding, rather than simply learning to perform an NLP task (Hewitt and Liang, 2019; Zhang and Bowman, 2018; Voita and Titov, 2020). This preference for simplicity has often led researchers to place restrictions on probe designs that may not allow them to fully exploit the structure in which information is encoded (Saphra and Lopez, 2019; Pimentel et al.,

2020b,a). This preference has led many researchers to advocate the use of linear probes over non-linear ones (Alain and Bengio, 2017).

This paper treats and expands upon the structural probe of Hewitt and Manning (2019), who crafted a custom probe with the aim of investigating the encoding of sentence syntax in contextual representations. They treat probing for syntax as a distance learning problem: they learn a linear transformation that warps the space such that two words that are syntactically close to one another (in terms of distance in a dependency tree) should have contextual representations whose Euclidean distance is small. This linear approach performs well, but the restriction to learning only linear transformations seems arbitrary. Why should it be the case that this information would be encoded linearly within the representations?

In this paper, we recast Hewitt and Manning’s (2019) structural probing framework as a general metric learning problem. This reduction allows us to take advantage of a wide variety of non-linear extensions—based on kernelization—proposed in the metric learning literature (Kulis, 2013). These extensions lead to probes with the *same* number of parameters, but with an increased expressivity.

By exploiting a kernelized extension, we are able to directly test whether a structural probe that is capable of learning non-linear transformations improves performance. Empirically, we do find that non-linearity helps—a structural probe based on a radial-basis function (RBF) kernel improves performance significantly in all 6 languages tested over a linear structural probe. We then perform an analysis of BERT’s attention, asserting it is a rough approximation to an RBF kernel. As such, it is not surprising that the syntactic information in BERT representations is more accessible with this specific *non-linear* transformation. We conclude that kernelization is a useful tool for analyzing contextual representations—enabling us to run controlled

experiments and investigate the structure in which information is encoded.

2 The Structural Probe

Hewitt and Manning (2019) introduce the **structural probe**, a novel model designed to probe for syntax in contextual word representations. We review their formulation here and build upon it in §4. A sentence \mathbf{w} lives in a space V^* , defined here as the Kleene closure of a (potentially open) vocabulary V . The syntactic distance Δ_{ij} between any two words in a sentence \mathbf{w} is the number of steps needed to go from one word to the other while walking in the sentence’s syntactic tree. More formally, if we have a dependency tree \mathbf{t} (a tree on $n + 1$ nodes) of a sentence \mathbf{w} of length n , we define Δ_{ij} as the length of the shortest path in \mathbf{t} between w_i and w_j ; this may be computed, for example, by Floyd–Warshall. Contextual representations of a sentence \mathbf{w} are a sequence of vectors $\mathbf{h}_i \in \mathbb{R}^{d_1}$ that encode some linguistic knowledge about a sequence. In the case of BERT, we have

$$\mathbf{h}_i = \text{BERT}(\mathbf{w})_i \in \mathbb{R}^{d_1} \quad (1)$$

Here, the goal of probing is to evaluate whether the contextual representations capture the syntax in a sentence. In the case of the structural probe, the goal is to see whether the syntactic distance between any two words can be approximated by a learned, linear distance function:

$$d_{\mathbf{B}}(\mathbf{h}_i, \mathbf{h}_j) = \|\mathbf{B}\mathbf{h}_i - \mathbf{B}\mathbf{h}_j\|_2 \quad (2)$$

where $\mathbf{B} \in \mathbb{R}^{d_2 \times d_1}$ is a linear projection matrix. That is to say, they seek a linear transformation such that the transformed contextual representations relate to one another roughly as their corresponding words do in the dependency tree. To learn this probe, Hewitt and Manning minimize the following per-sentence objective with respect to \mathbf{B} through stochastic gradient descent

$$\frac{1}{|\mathbf{w}|^2} \sum_{i=1}^{|\mathbf{w}|} \sum_{j=i+1}^{|\mathbf{w}|} |\Delta_{ij} - d_{\mathbf{B}}(\mathbf{h}_i, \mathbf{h}_j)| \quad (3)$$

This is simply minimizing the difference between the syntactic distances obtained from the dependency tree and the distance between the two vectors under our learned transformation. From the pairwise distances predicted by the probe, Prim’s (1957) algorithm can be used to recover the one-best undirected dependency tree.

3 Kernelized Metric Learning

The restriction to a linear transformation may hinder us from uncovering some of the syntactic structure encoded in the contextual representations. Indeed, there is no reason *a-priori* to expect that BERT encodes its knowledge in a fashion that is specifically accessible to a *linear* model. However, if we were to introduce non-linearity by using a neural probe, for example, we would have to pit a model with very few parameters (the linear model) against one with very many (the neural network); this comparison is not fair and also goes against the spirit of designing simple probes. To preclude the need for a neural probe, we instead turn to a kernelized probe.

The key insight is that the structural probe reduces the problem of probing for linguistic structure to that of **metric learning** (Kulis, 2013). This can be clearly seen in eq. (3), where the probe learns a distance metric between two representations in such a way that it matches the syntactic one. Recognizing this relationship allows us to take advantage of established techniques from the metric learning literature to improve the performance of the probe without increasing its complexity, e.g. through kernelization.

3.1 The “Kernel Trick” for Distances

Many algorithms in machine learning, e.g. support vector machines and k -means, can be **kernelized** (Schölkopf and Smola, 2002), thus allowing for linear models to be adapted into non-linear ones. Expanding on a classic result (Schoenberg, 1938), Schölkopf (2001) show that any positive semi-definite (PSD) kernel can be used to construct a distance in a Hilbert space \mathcal{H} . Formally, their result states that for any PSD kernel $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$, there exists a feature map $\phi : \mathcal{X} \rightarrow \mathcal{H}$ such that

$$\|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_2 = \sqrt{\kappa(\mathbf{x}, \mathbf{x}) - 2\kappa(\mathbf{x}, \mathbf{y}) + \kappa(\mathbf{y}, \mathbf{y})} \quad (4)$$

This generalizes eq. (2) to yield a new, non-linear distance metric. This means that we can achieve the effects of using some non-linear feature mapping ϕ without having to specify it: we need only specify a kernel function and perform calculations using this kernelized distance metric. Importantly, as opposed to deep neural probes, this learnable metric has an *identical* number of parameters to the original.¹

¹We note that we do not use selectivity (Hewitt and Liang, 2019) to control for probe complexity since it does not apply to

3.2 Common Kernels

In this section we introduce the kernels to be used. These kernels were chosen as they represent a comprehensive selection of commonly-used kernels in the metric learning literature (Kulis, 2013). The original work of Hewitt and Manning (2019) makes use of the **linear kernel**:

$$\kappa_{\text{linear}}(\mathbf{h}_i, \mathbf{h}_j) = (\mathbf{B}\mathbf{h}_i)^\top (\mathbf{B}\mathbf{h}_j) \quad (5)$$

The first non-linear kernel we consider is the **polynomial kernel**, defined as

$$\kappa_{\text{poly}}(\mathbf{h}_i, \mathbf{h}_j) = \left((\mathbf{B}\mathbf{h}_i)^\top (\mathbf{B}\mathbf{h}_j) + c \right)^d \quad (6)$$

where $d \in \mathbb{Z}_+$ and $c \in \mathbb{R}_{\geq 0}$. A polynomial kernel of degree d allows for d -order interactions between the terms. When working with BERT, this means that we may construct d -order conjunctions of the dimensions of the contextual representations input into the probe. Next, we consider the **radial-basis function kernel** (RBF). This kernel is also called the Gaussian kernel and is defined as

$$\kappa_{\text{rbf}}(\mathbf{h}_i, \mathbf{h}_j) = \exp\left(-\frac{\|\mathbf{B}\mathbf{h}_i - \mathbf{B}\mathbf{h}_j\|^2}{2\sigma^2}\right) \quad (7)$$

This kernel has an alternative interpretation as a similarity measure between both vectors, being at its maximum value of 1 when $\mathbf{h}_i = \mathbf{h}_j$. In contrast to the polynomial kernel, the Gaussian kernel implies a feature map in an infinite dimensional Hilbert space. When the RBF kernel is used in our probe, we may rewrite eq. (2) as follows:

$$\begin{aligned} d_{\kappa_{\text{rbf}}}(\mathbf{h}_i, \mathbf{h}_j)^2 & \quad (8) \\ &= \kappa_{\text{rbf}}(\mathbf{h}_i, \mathbf{h}_i) - 2\kappa_{\text{rbf}}(\mathbf{h}_i, \mathbf{h}_j) + \kappa_{\text{rbf}}(\mathbf{h}_j, \mathbf{h}_j) \\ &= 2 - 2\kappa_{\text{rbf}}(\mathbf{h}_i, \mathbf{h}_j) \\ &= 2 - 2 \exp\left(-\frac{\|\mathbf{B}\mathbf{h}_i - \mathbf{B}\mathbf{h}_j\|^2}{2\sigma^2}\right) \end{aligned}$$

Which is similar to the original linear case in eq. (2), but with a scaling term $-\frac{1}{2\sigma^2}$ and a non-linearity $\exp(\cdot)$. Finally, we consider, the **sigmoid kernel**, which is defined as ²

$$\kappa_{\text{sig}}(\mathbf{h}_i, \mathbf{h}_j) = \tanh(a(\mathbf{B}\mathbf{h}_i)^\top (\mathbf{B}\mathbf{h}_j) + b) \quad (9)$$

this syntax tree reconstruction task—selectivity control tasks work at the word type level, as opposed to the sentence one.

²Lin and Lin (2003) observe that it is difficult to effectively tune a and b in the sigmoid kernel. They also note that although this kernel is not in fact PSD, it is PSD when a and b are both positive, which we enforce in this work.

4 Regularized Metric Learning

We also take advantage of two common regularization techniques employed in the metric learning literature to further improve the transformations learned; both act on the matrix $\mathbf{A} = \mathbf{B}^\top \mathbf{B}$ and are added to the objective specified in eq. (3). The **Frobenius norm regularizer** takes the form

$$r(\mathbf{A}) = \|\mathbf{A}\|_{\text{F}}^2 = \text{tr}\left(\mathbf{A}^\top \mathbf{A}\right) \quad (10)$$

This is the matrix analogue of the L_2 squared regularizer. Minimizing the Frobenius norm of the learned matrix has the effect of keeping the values in the matrix small. It has been a popular choice for regularization in metric learning with adaptations to a variety of problems (Schultz and Joachims, 2004; Kwok and Tsang, 2003). We also consider the **trace norm regularizer**, which is of the form

$$r(\mathbf{A}) = \text{tr}(\mathbf{A}) \quad (11)$$

The trace norm regularizer is the matrix analogue of the L_1 regularizer and it encourages the matrix \mathbf{A} to be low rank. As Jain et al. (2010) point out, using a low-rank transformation in conjunction with a kernel corresponds to a supervised kernel dimensionality reduction method.

5 Experiments

We experiment with Hewitt and Manning’s (2019) probe on 6 typologically diverse languages, following the experimental design of Hall Maudslay et al. (2020). Our data comes from the Universal Dependency 2.4 Treebank (Nivre et al., 2019), providing sentences and their dependency trees, annotated using the Universal Dependencies annotation scheme.³ For each sentence we calculate contextual representations using multilingual BERT. For all languages, we took the first 12,000 sentences (or the maximum number thereof) in the train portion of the treebank and created new 80–10–10 train–test–dev splits.⁴

³It was recently demonstrated by Kuznetsov and Gurevych (2020) that choice of linguistic formalism may have an impact on probing results. In this work, we investigate using only one formalism, so we cannot be sure that our results would not differ if an alternative formalism were used. Nonetheless, we believe that the results that we find most interesting, which are discussed in §6, should be robust to a change in formalism, since their explanation lies in the way attention is calculated in the transformer architecture.

⁴We cap the maximum number of sentences analyzed as a naïve control for our multilingual analysis.

Kernel	Basque		English		Finnish		Korean		Tamil		Turkish	
	UUAS	DSpr	UUAS	DSpr	UUAS	DSpr	UUAS	DSpr	UUAS	DSpr	UUAS	DSpr
None	58.39	0.6737	57.96	0.7382	59.90	0.7560	68.63	0.7026	48.52	0.5116	58.87	0.6784
Polynomial	50.10	0.5751	59.67	0.7635	57.12	0.7401	67.58	0.6966	54.43	0.5776	55.29	0.6421
Sigmoid	43.14	0.4500	41.62	0.6152	53.14	0.6201	44.48	0.3734	44.30	0.3836	51.77	0.5557
RBF	60.99	0.6937	62.77	0.7213	63.08	0.7382	71.87	0.6918	56.96	0.5379	61.67	0.6841

Table 1: Results of probes using various kernels, in terms of UUAS and DSpr

We present the results from our comparison of a re-implementation of Hewitt and Manning’s (2019) linear structural probe and the non-linear kernelized probes in Table 1. The two evaluation metrics shown are unlabeled undirected attachment score (UUAS) and the Spearman rank-order correlation (DSpr) between predicted distances and gold standard pairwise distances. UUAS is a standard parsing metric expressing the percentage of correct attachments in the dependency tree, while DSpr is a measure of how accurately the probe predicts the overall ordering of distances between words. We can see that the use of an RBF kernel results in a statistically significant improvement in performance, as measured by UUAS, in all 6 of the languages tested.⁵ For some languages this improvement is quite substantial, with Tamil seeing an improvement of 8.44 UUAS from the baseline probe to the RBF kernel probe.

6 The RBF Kernel and Self-Attention

The RBF kernel produces improvements across all analyzed languages. This suggests that it is indeed the case that syntactic structure is encoded non-linearly in BERT. As such, analyzing this specific kernel may yield insights into what this structure is. Indeed, none of the other kernels systematically improve over the linear baseline, implying this is not just an effect of the non-linearity introduced through use of a kernel—the specific structure of the RBF kernel must be responsible. In this section, we argue that the reason that the RBF kernel serves as such a boon to probing is that it resembles BERT’s attention mechanism; recall that BERT’s attention mechanism is defined as

$$\text{att}(\mathbf{h}_i, \mathbf{h}_j) \propto \exp\left(\frac{(\mathbf{K}\mathbf{h}_i)^\top(\mathbf{Q}\mathbf{h}_j)}{\sqrt{d_2}}\right) \quad (12)$$

where \mathbf{K} and \mathbf{Q} are linear transformations and d_2 is the dimension vectors are projected into.

⁵Significance was established using paired permutation tests with 10,000 samples, to the level of $p < 0.05$.

\mathbf{K} projects vector \mathbf{h}_i into a key vector, while \mathbf{Q} projects \mathbf{h}_j into a query one. When the key and query vectors are similar (i.e. have a high dot product), the value of this equation is large and word j attends to word i .

This bears a striking resemblance to the Gaussian kernel. Indeed, if we assume the linearly transformed representations have unit norm, i.e.

$$\|\mathbf{B}\mathbf{h}_i\|^2 = \|\mathbf{B}\mathbf{h}_j\|^2 = 1 \quad (13)$$

then we have

$$\begin{aligned} & \exp\left(\frac{-1}{2\sqrt{d_2}}\|\mathbf{B}\mathbf{h}_i - \mathbf{B}\mathbf{h}_j\|^2\right) \\ &= \exp\left(\frac{-1}{\sqrt{d_2}} + \frac{(\mathbf{B}\mathbf{h}_i)^\top(\mathbf{B}\mathbf{h}_j)}{\sqrt{d_2}}\right) \\ &\propto \exp\left(\frac{(\mathbf{B}\mathbf{h}_i)^\top(\mathbf{B}\mathbf{h}_j)}{\sqrt{d_2}}\right) \end{aligned} \quad (14)$$

where we take $\sigma^2 = \sqrt{d_2}$. The similarity between eqs. (12) and (14) suggests the attention mechanism in BERT is, up to a multiplicative factor, roughly equivalent to an RBF kernel—as such, it is not surprising that the RBF kernel produces the strongest results.

The resemblance between these equations, taken together with the significant improvements in capturing syntactic distance, suggest that this encoded information indeed lives in an RBF-like space in BERT. Such information can then be used in its self-attention mechanism; allowing BERT to pay attention to syntactically close words when solving the cloze language modeling task. Being attentive to syntactically close words would also be supported by recent linguistic research, since words sharing syntactic dependencies have higher mutual information on average (Futrell et al., 2019).

The representations we analyze, though, are taken from BERT’s final layer; as such, they are not trained to be used in any self-attention layer—so why should such a resemblance be relevant? BERT’s architecture is based on the Transformer

(Vaswani et al., 2017), and uses skip connections between each self-attention layer. Such skip connections create an incentive for residual learning, i.e. only learning residual differences in each layer, while propagating the bulk of the information (He et al., 2016). As such, BERT’s final hidden representations should roughly live in the same manifold as its internal ones.

It is interesting to note that the RBF kernel achieves the best performance in terms of UUAS in all languages, but it only twice achieves the best performance in terms of DSpr. This may be due to the fact that, as we can see by examination of eq. (8), the distance returned by the RBF kernel will not exceed 2, whereas syntactic distances in the tree will. Further, the gradient of the RBF kernel contains an exponential term which will cause it to go to zero as distance increases (while an examination of the unkernelized loss function reveals the opposite behavior). This means that it will be less sensitive to the distances between syntactically distant words and focus more on words with small distances. This may partially explain its better performance on UUAS, and comparably worse performance as measured by correlation (which counts pairwise differences between *all* words, not just those which are directly attached in the tree). Furthermore, our probe’s focus on nearby words resembles the general attentional bias towards syntactically close words (Voita et al., 2019).

The direct resemblance between self-attention mechanisms and our proposed probe metric poses a new way of understanding results from more complex probes. While Reif et al. (2019) understood the Euclidean-squared distance of Hewitt and Manning as an isometric tree embedding, their geometric interpretation did not factor in the rest of BERT’s architecture. Such simplified contextless probes cannot tell us how linguistic properties are processed by a sequence of learned modules (Saphra and Lopez, 2019). However, we consider representations *in the context* of the model which is expected to employ them. From this perspective, simpler metrics may be rough approximations to our RBF kernel space, which is actually capable of measuring linguistic properties that can be easily extracted by an attention-based architecture.

7 Conclusion

We find that the linear structural probe (Hewitt and Manning, 2019) used to investigate the encoding

of syntactic structure in contextual representations can be improved through kernelization, yielding a non-linear model. This kernelization does not introduce additional parameters and thus does not increase the complexity of the probe—at least if one treats the number of parameters as a good proxy for model complexity. At the same time, the RBF kernel improves probe performance in all languages under consideration. This suggests that syntactic information may be encoded non-linearly in the representations produced by BERT. We hypothesize that this is true due to the similarity of the RBF kernel and BERT’s self-attention layers.

Ethical Considerations

The authors foresee no ethical concerns with the research presented in this paper.

References

- Guillaume Alain and Yoshua Bengio. 2017. [Understanding intermediate layers using linear classifier probes](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Robert W. Floyd. 1962. Algorithm 97: Shortest path. *Communications of the ACM*, 5(6):345.
- Richard Futrell, Peng Qian, Edward Gibson, Evelina Fedorenko, and Idan Blank. 2019. [Syntactic dependencies correspond to word pairs with high mutual information](#). In *Proceedings of the Fifth International Conference on Dependency Linguistics (Depling, SyntaxFest 2019)*, pages 3–13, Paris, France. Association for Computational Linguistics.
- Rowan Hall Maudslay, Josef Valvoda, Tiago Pimentel, Adina Williams, and Ryan Cotterell. 2020. [A tale of a probe and a parser](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.

- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Prateek Jain, Brian Kulis, and Inderjit Dhillon. 2010. [Inductive regularized learning of kernel functions](#). In *Advances in Neural Information Processing Systems*, volume 23.
- Brian Kulis. 2013. *Metric Learning: A Survey*, volume 5 of *Foundations and Trends® in Machine Learning*. Now Publishers, Inc.
- Iliia Kuznetsov and Iryna Gurevych. 2020. [A matter of framing: The impact of linguistic formalism on probing results](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 171–182, Online. Association for Computational Linguistics.
- James T. Kwok and Ivor W. Tsang. 2003. [Learning with idealized kernels](#). In *Proceedings of the 20th International Conference on Machine Learning*, pages 400–407.
- Hsuan-Tien Lin and Chih-Jen Lin. 2003. A study on sigmoid kernels for SVM and the training of non-PSD kernels by SMO-type methods. *Neural Computation*, 3:1–32.
- Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Gabrielė Aleksandravičiūtė, Lene Antonsen, Katya Aplonova, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Victoria Basmov, John Bauer, Sandra Bellato, Kepa Bengoetxea, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Agnė Bielinskienė, Rogier Blokland, Victoria Bobicev, Loïc Boizou, Emanuel Borges Völker, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Kristina Brokaitė, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čéplö, Savas Cetin, Fabricio Chalub, Jinho Choi, Yongseok Cho, Jayeol Chun, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Arantza Diaz de Ilaraza, Carly Dickerson, Bamba Dione, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Hanne Eckhoff, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Tomaž Erjavec, Aline Etienne, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Kazunori Fujita, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Matias Grioni, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Nizar Habash, Jan Hajič, Jan Hajič jr., Linh Hà Mỷ, Na-Rae Han, Kim Harris, Dag Haug, Johannes Heinecke, Felix Hennig, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Takumi Ikeda, Radu Ion, Elena Irimia, Olájídé Ishola, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Hüner Kaşıkara, Andre Kaasen, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Arne Köhn, Kamil Kopacewicz, Natalia Kotsyba, Jolanta Kovalevskaitė, Simon Krek, Sookyoung Kwak, Veronika Laippala, Lorenzo Lambertino, Lucia Lam, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phương Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, Kyung-Tae Lim, Yuan Li, Nikola Ljubešić, Olga Loginova, Olga Lyashevskaya, Teresa Lynn, Vivien Mackentanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Sarah McGuinness, Gustavo Mendonça, Niko Miekka, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Keiko Sophie Mori, Tomohiko Morioka, Shinsuke Mori, Shigeki Moro, Bjartur Mortensen, Bohdan Moskalevskiy, Kadri Muischnek, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horňiáček, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Lương Nguyễn Thị, Huyền Nguyễn Thị Minh, Yoshihiro Nikaïdo, Vitaly Nikolaev, Rattima Nitisaroj, Hanna Nurmi, Stina Ojala, Adédayo Olúòkun, Mai Omura, Petya Osenova, Robert Östling, Lilja Øvrelid, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Angelika Peljak-Łapińska, Siyao Peng, Ceneľ-Augusto Perez, Guy Perrier, Daria Petrova, Slav Petrov, Jussi Piitulainen, Tommi A Pirinen, Emily Pitler, Barbara Plank, Thierry Poibeau, Martin Popel, Lauma Pretkalnina, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Michael Riebler, Erika Rimkutė, Larissa Rinaldi, Laura Rituma, Luisa

- Rocha, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Valentin Roșca, Olga Rudina, Jack Rueter, Shoval Sadde, Benoît Sagot, Shadi Saleh, Alessio Salomoni, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Dage Särög, Baiba Saulīte, Yanin Sawanakunanon, Nathan Schneider, Sebastian Schuster, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Hiroyuki Shirasu, Muh Shohibussirri, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Isabela Soares-Bastos, Carolyn Spadine, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Shingo Suzuki, Zsolt Szántó, Dima Taji, Yuta Takahashi, Fabio Tamburini, Takaaki Tanaka, Isabelle Tellier, Guillaume Thomas, Lisi Torga, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Urešová, Larraitz Uribe, Hans Uszkoreit, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villemonte de la Clergerie, Veronika Vincze, Lars Wallin, Abigail Walsh, Jing Xian Wang, Jonathan North Washington, Maximilian Wendt, Seyi Williams, Mats Wirén, Christian Wittern, Tsegay Woldemariam, Tak-sum Wong, Alina Wróblewska, Mary Yako, Naoki Yamazaki, Chunxiao Yan, Koichi Yasuoka, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, Manying Zhang, and Hanzhi Zhu. 2019. [Universal Dependencies 2.4](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Matthew 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)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Tiago Pimentel, Naomi Saphra, Adina Williams, and Ryan Cotterell. 2020a. [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.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020b. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Robert Clay Prim. 1957. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019. [Visualizing and measuring the geometry of bert](#). In *Advances in Neural Information Processing Systems*, volume 32.
- Naomi Saphra and Adam Lopez. 2019. [Understanding learning dynamics of language models with SVCCA](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3257–3267, Minneapolis, Minnesota. Association for Computational Linguistics.
- I. J. Schoenberg. 1938. Metric spaces and positive definite functions. *Transactions of the American Mathematical Society*, 44(3):522–536.
- Bernhard Schölkopf. 2001. [The kernel trick for distances](#). In *Advances in Neural Information Processing Systems*, pages 301–307.
- Bernhard Schölkopf and Alexander J. Smola. 2002. *Learning with Kernels*. MIT Press.
- Matthew Schultz and Thorsten Joachims. 2004. [Learning a distance metric from relative comparisons](#). In *Advances in Neural Information Processing Systems*, pages 41–48.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Elena Voita and Ivan Titov. 2020. [Information-theoretic probing with minimum description length](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 183–196, Online. Association for Computational Linguistics.
- Stephen Warshall. 1962. A theorem on boolean matrices. *Journal of the ACM*, 9(1):11–12.
- 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.