

The Rediscovery Hypothesis: Language Models Need to Meet Linguistics

Vassilina Nikoulina

NAVER LABS Europe

6-8 chemin de Maupertuis, 38240 Meylan, France

VASSILINA.NIKOULINA@NAVERLABS.COM

Maxat Tezekbayev

Nuradil Kozhakhmet

Madina Babazhanova

Nazarbayev University

53 Kabanbay Batyr ave., Nur-Sultan Z05H0P9, Kazakhstan

MAXAT.TEZEKBAYEV@NU.EDU.KZ

NURADIL.KOZHAKHMET@NU.EDU.KZ

MADINA.BABAZHANOVA@NU.EDU.KZ

Matthias Gallé

NAVER LABS Europe

6-8 chemin de Maupertuis, 38240 Meylan, France

MATTHIAS.GALLE@NAVERLABS.COM

Zhenisbek Assylbekov

Nazarbayev University

53 Kabanbay Batyr ave., Nur-Sultan Z05H0P9, Kazakhstan

ZHASSYLBEKOV@NU.EDU.KZ

Abstract

There is an ongoing debate in the NLP community whether modern language models contain linguistic knowledge, recovered through so-called *probes*. In this paper, we study whether linguistic knowledge is a necessary condition for the good performance of modern language models, which we call the *rediscovery hypothesis*.

In the first place, we show that language models that are significantly compressed but perform well on their pretraining objectives retain good scores when probed for linguistic structures. This result supports the rediscovery hypothesis and leads to the second contribution of our paper: an information-theoretic framework that relates language modeling objectives with linguistic information. This framework also provides a metric to measure the impact of linguistic information on the word prediction task. We reinforce our analytical results with various experiments, both on synthetic and on real NLP tasks in English.

1. Introduction

Vector representations of words obtained from self-supervised pretraining of neural language models (LMs) on massive unlabeled data have revolutionized NLP in the last decade. This success has spurred an interest in trying to understand what type of knowledge these models actually learn (Rogers et al., 2020; Petroni et al., 2019a).

Of particular interest here is “linguistic knowledge”, which is generally measured by annotating test sets through experts (linguists) following certain pre-defined linguistic schemas. These annotation schemas are based on language regularities manually defined by linguists. On the other side, we have language models which are pretrained predictors that assign a probability to the presence of a token given the surrounding context. Neural models solve this task by finding patterns in the text. We refer to the claim that such neural language models rediscover linguistic knowledge as the *rediscovery hypothesis*. It stipulates that the

patterns of the language discovered by the model trying to solve the pretraining task correlate with the human-defined linguistic regularity. In this work we measure the amount of linguistics rediscovered by a pretrained model through the so-called probing tasks: hidden layers of a neural LM are fed to a simple classifier—a probe—that learns to predict a linguistic structure of interest (Ettinger et al., 2016; Adi et al., 2017; Conneau et al., 2018; Hewitt & Manning, 2019; Tenney et al., 2019a). In the first part of this paper (Section 2) we attempt to challenge the *rediscovery hypothesis* through a variety of experiments to understand to what extent it holds. Those experiments aim to verify whether the path through language regularity is indeed the one taken by pretrained LMs or whether there is another way to reach good LM performance without rediscovering linguistic structure. Our experiments show that pretraining loss is indeed tightly linked to the amount of linguistic structure discovered by an LM. We, therefore, fail to reject the rediscovery hypothesis.

This negative attempt, as well as the abundance of positive examples in the literature, motivates us to prove mathematically the rediscovery hypothesis. In the second part of our paper (Section 3) we use information theory to prove the contrapositive of the hypothesis—removal of linguistic information from an LM degrades its performance. Moreover, we show that the decline in the LM quality depends on how strongly the removed property is interdependent with the underlying text: a greater dependence leads to a greater drop. We confirm this result empirically, both with synthetic data and with real annotations on English text.

The result that removing information that contains strong mutual information with the underlying text degrades (masked) word prediction might not seem surprising *a posteriori*. However, it is this surprise that lies at the heart of most of the work in recent years around the discovery of how easily this information can be extracted from intermediate representations. Our framework also provides a coefficient that measures the dependence between a probing task and the underlying text. This measure can be used to determine more complex probing tasks, whose rediscovery by language models would indeed be surprising.

2. Do Pretraining Objectives Correlate with Linguistic Structures?

The first question we pose in this work is the following: is the rediscovery of linguistic knowledge mandatory for models that perform well on their pretraining tasks, typically language modeling or translation; or is it a side effect of overparameterization?¹

We analyze the correlation between linguistic knowledge and LM performance with pruned pretrained models. By compressing a network through pruning we retain the same overall architecture and can compare probing methods. More important, we hypothesize that pruning removes all unnecessary information with respect to the pruning objective (language modeling) and that it might be that information that is used to rediscover linguistic knowledge. To complete this experiment, we also track the pruning efficiency *during* pretraining.

1. Overparameterization is defined informally as “having more parameters than can be estimated from the data”, and therefore using a model richer than necessary for the task at hand. Those additional parameters could be responsible for the good performance of the probes.

2.1 Pruning Method

The lottery ticket hypothesis (LTH) of Frankle and Carbin (2019) claims that a randomly-initialized neural network $f(\mathbf{x}; \boldsymbol{\theta})$ with trainable parameters $\boldsymbol{\theta} \in \mathbb{R}^n$ contains subnetworks $f(\mathbf{x}; \mathbf{m} \odot \boldsymbol{\theta})$, $\mathbf{m} \in \{0, 1\}^n$, such that, when trained in isolation, they can match the performance of the original network.² The authors suggest a simple procedure for identifying such subnetworks (Alg. 2.1). When this procedure is applied iteratively (Step 5 of Alg. 2.1), we get a sequence of pruned models $\{f(\mathbf{x}; \mathbf{m}_i \odot \boldsymbol{\theta}_0)\}$ in which each model has fewer parameters than its predecessor: $\|\mathbf{m}_i\|_0 < \|\mathbf{m}_{i-1}\|_0$. Frankle and Carbin used iterative pruning for image classification models and found subnetworks that were 10%–20% of the sizes of the original networks and met or exceeded their validation accuracies. Such a compression approach retains weights important for the main task while discarding others. We hypothesize that it might be those additional weights that contain the signals used by probes. But, if the subnetworks retain linguistic knowledge then this is evidence in favor of the rediscovery hypothesis.

Algorithm 2.1: Lottery ticket hypothesis—Identifying winning tickets (Frankle & Carbin, 2019)

- 1 Randomly initialize a neural network $f(\mathbf{x}; \boldsymbol{\theta}_0)$, $\boldsymbol{\theta}_0 \in \mathbb{R}^n$
 - 2 Train the network for j iterations, arriving at parameters $\boldsymbol{\theta}_j$.
 - 3 Prune $p\%$ of the parameters in $\boldsymbol{\theta}_j$, creating a mask $\mathbf{m} \in \{0, 1\}^n$.
 - 4 Reset the remaining parameters to their values in $\boldsymbol{\theta}_0$, creating the winning ticket $f(\mathbf{x}; \mathbf{m} \odot \boldsymbol{\theta}_0)$.
 - 5 Repeat from 2 if performing iterative pruning.
 - 6 Train the winning ticket $f(\mathbf{x}; \mathbf{m} \odot \boldsymbol{\theta}_0)$ to convergence.
-

2.2 Models

We explore one static embedding model, SGNS, and two contextualized embedding models, CoVe and RoBERTa. In this manner, we cover the full spectrum of modern representations of words, from static embeddings to shallow and deep contextual embeddings.

CoVe (McCann et al., 2017) uses the top-level activations of a two-layer BiLSTM encoder from an attentional sequence-to-sequence model (Bahdanau et al., 2015) trained for English-to-German translation. The authors used the CommonCrawl-840B GLoVe model (Pennington et al., 2014) for English word vectors, which were completely fixed during pretraining, and we follow their setup. This entails that the embedding layer on the source side is not pruned during the LTH procedure. We also concatenate the encoder output with the GLoVe embeddings as is done in the original paper (McCann et al., 2017, Eq. 6).

BERT (Devlin et al., 2019) is a deep Transformer (Vaswani et al., 2017) encoder that has become the *de facto* standard when it comes to contextualized embeddings. We pretrain the RoBERTa variant of the BERT model (Liu et al., 2019). RoBERTa stands for robustly optimized BERT which was trained with hyperparameters optimized for convergence, with

2. $\|\mathbf{x}\|_0$ is a number of nonzero elements in $\mathbf{x} \in \mathbb{R}^d$, \odot is the element-wise multiplication.

dynamic masking instead of static masking, and restricted to masked LM objective only.³ Unlike their predecessors, such as CoVe and ELMO (Peters et al., 2018), BERT and other Transformer-based encoders are considered deep contextualizers.

Word2vec SGNS (Mikolov et al., 2013b) is a shallow two-layer neural network that produces uncontextualized word embeddings. It is widely accepted that the SGNS vectors capture words semantics to a certain extent, which is confirmed by the folklore examples, such as $\mathbf{w}_{\text{king}} - \mathbf{w}_{\text{man}} + \mathbf{w}_{\text{woman}} \approx \mathbf{w}_{\text{queen}}$. So, the question of the relationship between the SGNS objective function and its ability to discover linguistics is also relevant.

2.3 Measuring the Amount of Linguistic Knowledge

To properly measure the *amount of linguistic knowledge* in word vectors we define it as the performance of classifiers (probes) that take those vectors as input and are trained on linguistically annotated data. This definition has the advantage of being able to be measured exactly, at the cost of avoiding the discussion of whether POS tags or syntactic parse trees indeed denote linguistic knowledge captured by humans in their learning process.

We should note that this probing approach has received a lot of criticism recently (Hewitt & Liang, 2019; Pimentel et al., 2020b; Voita & Titov, 2020) due to its inability to distinguish between information encoded in the pretrained vectors from the information learned by the probing classifier. However, in our study, the question is not *How much linguistics is encoded in the presentation vector?*, but rather *Does one vector contain more linguistic information than the other?* We compare different representations of the *same* dimensionality using probing classifiers of the *same* capacity. Even if part of the probing performance is due to the classifier itself we claim that the difference in the probing performance will be due to the difference in the amount of linguistic knowledge encoded in the representations we manipulate. This conjecture is strengthened by the findings of Zhang et al. (2020) who analyzed the representations from pretrained MINIBERTs⁴ and demonstrated that the trends found through edge probing (Tenney et al., 2019b) are the same as those found through better-designed probes such as Minimum Description Length (Voita & Titov, 2020). Therefore in our work we adopt *edge probing* and *structural probing* for contextualized embeddings. For static embeddings, we use the traditional *word similarity* and *word analogy* tasks.

Edge probing (Tenney et al., 2019b) formulates several linguistics tasks of different nature as text span classification tasks. The probing model is a lightweight classifier on top of the pretrained representations trained to solve those linguistic tasks. In our study, we use the part-of-speech tagging (POS), constituent labeling, named entity labeling (NE), and semantic role labeling (SRL) tasks from the suite, in which a probing classifier receives a sequence of tokens and predicts a label for it. For example, in the case of constituent labeling, for a sentence *This probe [discovers linguistic knowledge]*, the sequence in square brackets should be labeled as a verb phrase.

Structural probing (Hewitt & Manning, 2019) evaluates whether syntax trees are embedded in a linear transformation of a neural network’s word representation space. The

3. BERT has next sentence prediction loss in addition.

4. RoBERTa models trained on a varied amount of training data

probe identifies a linear transformation under which squared Euclidean distance encodes the distance between words in the parse tree. Hewitt and Manning show that such transformations exist for both ELMO and BERT but not in static baselines, providing evidence that entire syntax trees can be easily extracted from the vector geometry of deep models.

Word similarity (Finkelstein et al., 2002) and word analogy (Mikolov et al., 2013c) tasks can be considered as nonparametric probes of static embeddings, and—differently from the other probing tasks—are not learned. The use of word embeddings in the word similarity task has been criticized for the instability of the results obtained (Antoniak & Mimno, 2018). Regarding the word analogy task, Schlueter (2018) raised concerns on the misalignment of assumptions in generating and testing word embeddings. However, the success of the static embeddings in performing well in these tasks was a crucial part of their widespread adoption.

2.4 Experimental Setup

We prune the embedding models from Section 2.2 with the LTH algorithm (Alg. 2.1) and evaluate them with probes from Section 2.3 at *each* pruning iteration. COVE and SGNS are pruned iteratively, while for ROBERTA we perform one-shot pruning at different rates.⁵ Assuming that ℓ_i^ω is the validation loss of the embedding model $\omega \in \{\text{COVE}, \text{ROBERTA}, \text{SGNS}\}$ at iteration i , and $\Delta s_i^{\omega, T} := s_i^{\omega, T} - s_0^{\omega, T}$ is the drop in the corresponding score on the probing task $T \in \{\text{NE}, \text{POS}, \text{Const.}, \text{Struct.}, \text{Sim.}, \text{Analogy}\}$ compared to the score $s_0^{\omega, T}$ of the baseline (unpruned) model, we obtain pairs $(\ell_i^\omega, \Delta s_i^{\omega, T})$ for further analysis. Keep in mind that ROBERTA and SGNS are pruned in full, while in COVE we prune everything except the source-side embedding layer. This exception is due to the design of the COVE model, and we follow the original paper’s setup (McCann et al., 2017).

Software and datasets. We pretrain COVE on the English–German part of the IWSLT 2016 machine translation task (Cettolo et al., 2016) using the OPENNMT-PY toolkit (Klein et al., 2017). ROBERTA is pretrained on the WikiText-103 dataset (English) (Merity et al., 2017) using the FAIRSEQ toolkit (Ott et al., 2019b) with default training settings (Ott et al., 2019a). Finally, the SGNS model is pretrained on the `text8` data (English as well) (Mahoney, 2011) using our custom implementation (Assylbekov, 2020).

The edge probing classifier is trained on the standard benchmark dataset OntoNotes 5.0 (Weischedel et al., 2013) using the JIANT toolkit (Pruksachatkun et al., 2020). The structural probe is trained on the English UD (Silveira et al., 2014) using the code from the authors (Hewitt, 2019). For word similarities we use the WordSim353 dataset (Finkelstein et al., 2002), while for word analogies we use the Google dataset (Mikolov et al., 2013a). All those datasets are in English.

Optimization is performed in almost the same way as in the original works on COVE, ROBERTA, and SGNS. See Appendix A for details.

5. This is done to speedup experiments on ROBERTA, as one-shot pruning at different rates can be run in parallel.

2.5 Results

First, we note that the lottery ticket hypothesis is confirmed for the embedding models since pruning up to 60% weights does not harm their performance significantly on held-out data (Fig. 1).⁶

Before proceeding to the results of probing of pruned models, we want to note that the probing scores of the baseline (unpruned) models obtained by us are close to the scores from the original papers (Tenney et al., 2019b; Hewitt & Manning, 2019), and are shown in Table 1.

Model	Task				Model	Task	
	NE	POS	Const.	Struct.		Similarity	Analogy
CoVE	.921	.936	.808	.726	SGNS	.716	.332
ROBERTA	.932	.951	.808	.719			

Table 1: Probing scores for the baseline (unpruned) models. We report the micro-averaged F1 score for the POS, NE, and Constituents; undirected unlabeled attachment score (UAS) for the structural probe; Spearman’s correlation with the human ratings for the similarity task; and accuracy for the analogy task.

Probing results are provided in Fig. 2 and 3, where we scatter-plot validation loss ℓ_i^ω vs drop in probing performance $\Delta s_i^{\omega,T}$ for each of the model-probe combinations.⁷ First, we note that, in most cases, the probing score correlates with the pretraining loss, which supports the rediscovery hypothesis. We note that the probing score decreases slower for some tasks (e.g., POS tagging), but is much steeper for others (e.g., constituents labeling). This is complementary to the findings of Zhang et al. (2020) who showed that the syntactic learning curve reaches plateau performance with less pretraining data while solving semantic tasks requires more training data. Our results suggest that similar behavior emerges with respect to the model size: simpler tasks (e.g., POS tagging) can be solved with smaller models, while more complex linguistic tasks (e.g., syntactic constituents or dependency parsing) require bigger model size. In addition, we note that in the case of CoVE, and in contrast to ROBERTA, the probing scores for more local (i.e. less context-dependent) tasks such as POS and NER hardly decrease with an increase in the pruning rate. We believe that this is because CoVE representations by default contain unpruned static GLOVE embeddings, which by themselves already obtain good performance on the more local tasks.

The reader might have noted that in the case of CoVE in Fig. 2 when the range of ℓ_i ’s is restricted to low values, there is a lack of correlation between ℓ_i and Δs_i . We discuss this further in Appendix B as we argue that this does not contradict our main finding.

6. In the case of SGNS, pruning up to 80% of weights does not affect its validation loss. Since solving SGNS objective is essentially a factorization of the pointwise mutual information matrix, in the form $\text{PMI} - \log k \approx \mathbf{WC}$ (Levy & Goldberg, 2014), this means that a factorization with sparse \mathbf{W} and \mathbf{C} is possible. This observation complements the findings of Tissier et al. (2019) who showed that near-to-optimal factorization is possible with binary \mathbf{W} and \mathbf{C} .

7. Note, that cross-entropy values of CoVE and ROBERTA are not comparable as these are over different corpora, languages and vocabularies.

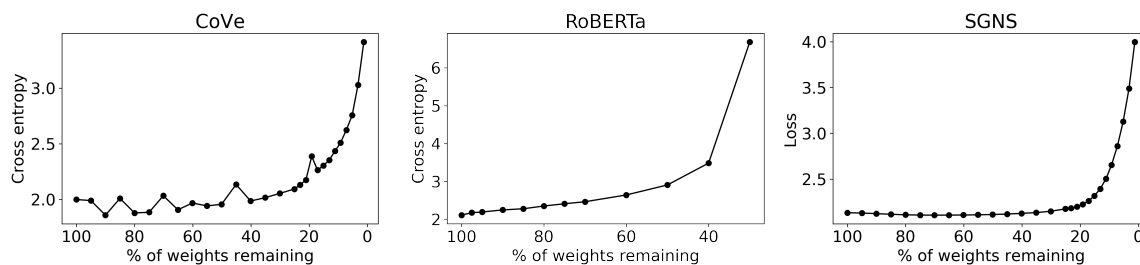


Figure 1: Results of applying the LTH (Algorithm 2.1) to CoVe, RoBERTa, and SGNS. In each case we scatter-plot the percentage of remained weights vs validation loss, which is cross entropy for CoVe and RoBERTa, and a variant of negative sampling objective for SGNS.




Figure 2: Probing results for the CoVe and RoBERTa embeddings. Horizontal axes indicate validation loss values, which are cross-entropy values. Vertical axes indicate drops in probing performances. In case of edge probing, we use the NE, POS, and constituent labeling tasks from the suite of Tenney et al. (2019b) and report the drop in micro-averaged F1 score compared to the baseline (unpruned) model. In case of structural probing, we use the distance probe of Hewitt and Manning (2019) and report the drop in undirected unlabeled attachment score (UJAS).

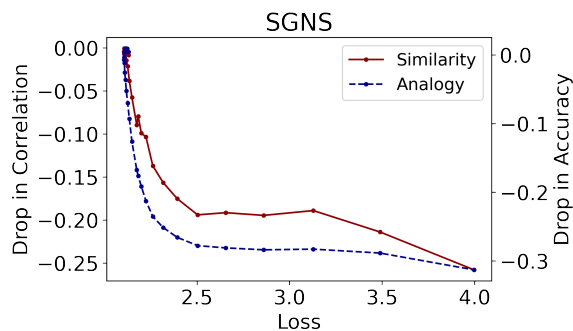


Figure 3: Similarity and analogy results for the SGNS embeddings. For similarities we report the drop in Spearman’s correlation with the human ratings and for analogies in accuracy.

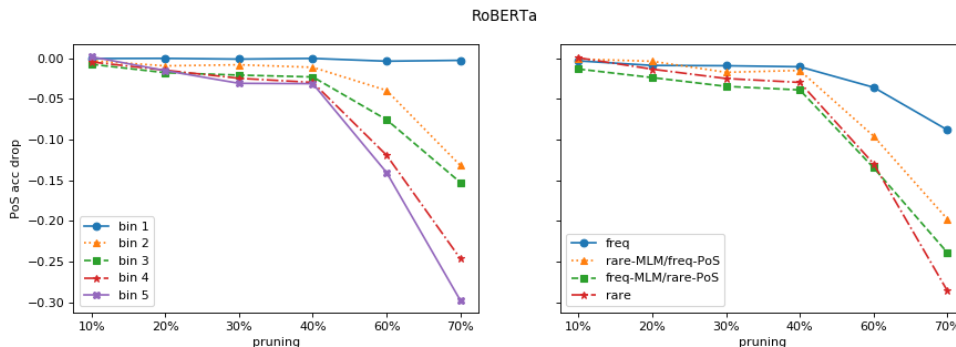


Figure 4: Breakdown of POS tagging accuracy decrease by token frequency. We report the drop in accuracy compared to the baseline (unpruned) model. **Left:** all bins (1-5) have comparable cumulative counts in the pretraining data. The distribution of tokens over the bins: bin 1 — 5 tokens, bin 2 — 25 tokens, bin 3 — 368 tokens, bin 4 — 2813 tokens, bin 5 — 600K tokens. **Right:** 4 bins of tokens which are (1) frequent both in pretraining and POS training data (153 tokens), (2) rare in pretraining but frequent in POS training data (66 tokens), (3) frequent in pretraining but rare in POS training data (235 tokens), and (4) rare everywhere (46K tokens).

Breakdown by token frequency. The pervasiveness of Zipf’s law in language has the consequence that many conclusions over aggregated scores can be attributed to effects on a small set of tokens. Indeed, when binning tokens by their frequency in the pretraining data and analyzing the POS accuracy drop per bin it becomes obvious that the less frequent a token is the more pruning affects its POS-tag prediction. This is shown in Fig. 4 (left) where tokens are binned in 5 equally-sized groups.

A straightforward interpretation of such behavior would be: (1) pruning degrades the representations for rarer tokens in the first place. It is also possible, that (2) pruning degrades all token representations similarly; however, the probing model has the capacity to recover the POS performance for tokens that are frequent in POS tagging corpus but cannot do that for rare tokens. The first reason would support the claim that pruning removes the *memorization*⁸ for rare tokens, while the second reason would mean that the pruned model redistributes its representation power across all the token groups to preserve good LM performance. The latter behavior would be more aligned with the *rediscovery hypothesis*.

To better understand which of the two is a better explanation we need to distinguish between tokens which are rare in the pretraining data but are frequent in the downstream POS training data and vice-versa: the behavior on tokens which are rare in POS training data will be more informative of how pruning affects pretrained representations.

8. To do well on the word predictions task the language model can either memorize patterns or learn certain “language regularity” (aka generalization). If we do an analogy with linguistic rules: there are cases when the linguistic rules (\sim generalization) can be applied, and there are exceptions. The generalization has its limits, and at some point, the language model needs to memorize to boost the performance.

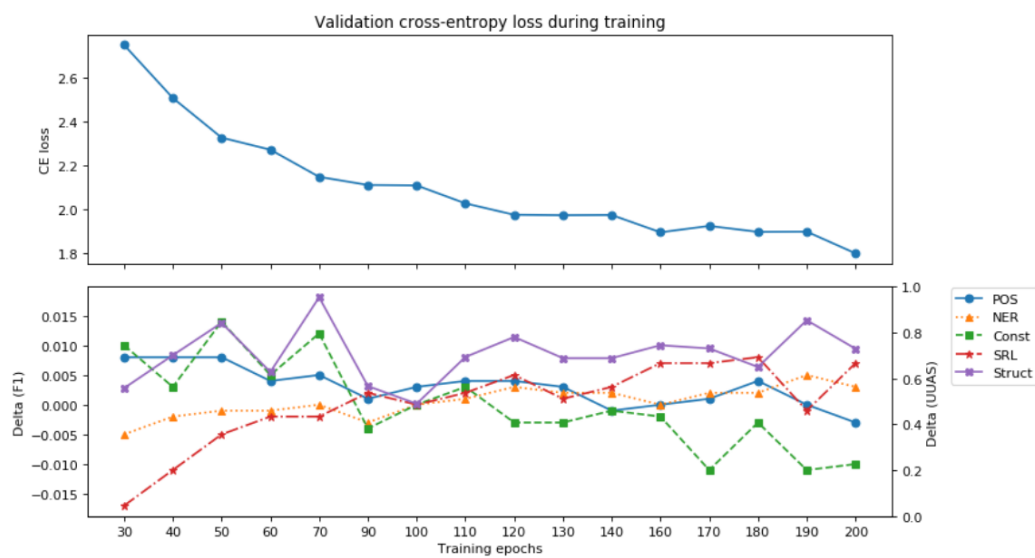


Figure 5: Probing intermediate checkpoints of RoBERTa. Top: RoBERTa validation CE loss during training. Bottom: Difference between the baseline probing score¹⁰ and probing score at each epoch (up to 200). In case of edge probing, we use the SRL, NE, POS, and constituent labeling tasks from the suite of Tenney et al. (2019b) and report the micro-averaged F1 score. In case of structural probing, we use the distance probe of Hewitt and Manning (2019) and report the undirected unlabeled attachment score (UUAS).

For this, we group the first three bins together in a **freq-MLM** group and the two last bins inside **rare-MLM**; and do a similar split based on the frequency of the tokens in the POS training corpus (**freq-POS** and **rare-POS**).

The impact of pruning on the four possible combinations of groups is shown in Fig. 4 (right). The probing performance on the tokens which are frequent in POS tagging corpus (**freq** and **rare-MLM/freq-POS**) keep almost constant up to a pruning rate of 40%. While tokens which are rare in POS tagging corpus (**freq-MLM/rare-POS** and **rare**) seem to suffer more with pruning rate growth. This lends support to the possibility that the amount of linguistics contained in the pretrained vectors decreases across all the token groups.

Finally, we note that up to 60% of pruning both **freq-MLM/rare-PoS** and **rare** groups behave similarly but at 70% pruning rate the tokens that are rare everywhere (**rare**) suffer from a higher drop, compared to **freq-MLM/rare-PoS** meaning that probing model is not able to recover the correct PoS tags for rare tokens. This suggests that *memorization* gets removed from the pretrained model at a higher pruning rate, but it is exploited by the probing model at lower pruning rates.

¹⁰. Baseline score corresponds to the probing with best checkpoint (around 450 epochs of training).

2.6 Probing Intermediate Checkpoints

Until now we measured correlated probing accuracy on top of fully-trained models with the amount of pruning those models underwent. Here we use a complementary lens, by analyzing how quickly a pretrained model discovers different types of linguistic knowledge? For this, we save intermediate checkpoints of RoBERTa models and probe those. Similar to pruning, this has the advantage that all models share the same architecture. The results of the probing accuracy over different epochs are provided in Fig. 5.

Similar to Zhang et al. (2020), we find that syntactic tasks (e.g., POS tagging, dependency, and constituency parsing) seem to reach the top probing performance at the beginning of training (30-60 epochs), while semantic tasks (e.g., NER, SRL) keep improving further. While Tenney et al. (2019a) argued that “*BERT rediscovers NLP pipeline*” by looking at intermediate layers of a fully pretrained model, the same seems to apply *during* the training: linguistic knowledge to solve lexical tasks are learned first while more complex tasks are easier solved with representation obtained later in the learning.

Inter-Section Interlude

We conclude that in general the rediscovery hypothesis seems valid: when language models reach good performance they indeed capture linguistic structures (at least in the case of the English language). So we obtained a negative result when we tried to reject the rediscovery hypothesis. Along with this, positive examples are abundant, including the extensive literature on probing (of which we give an overview in Section 4). This might indicate that the hypothesis is indeed true. The lack of correlation between probing scores and LM performance for one of the considered models only questions the probing methodology but does not reject the very fundamental connection between language modeling and learning linguistic structures. We address this in the next section, where we formalize the connection and show how it holds empirically. The empirical experiments will consist in adversarially removing information that could serve probing accuracy while keeping good LM performance.¹¹

3. An Information-Theoretic Framework

Recall that the rediscovery hypothesis asserts that neural language models, in the process of their pretraining, rediscover linguistic knowledge. We will prove this claim by contraposition, which states that without linguistic knowledge, the neural LMs cannot perform at their best. A recent paper of Elazar et al. (2021) has already investigated how linearly removing¹² certain linguistic information from BERT’s layers impacts its accuracy of predicting a masked token. They showed *empirically* that dependency information, part-of-speech tags, and named entity labels are important for word prediction, while syntactic constituency boundaries (which mark the beginning and the end of a phrase) are not. One of the questions raised by the authors is how to *quantify* the relative importance of different properties encoded in the representation for the word prediction task. The current section of our work

11. A more restricted scenario, where such adversarial training is performed on CoVE is reported in Appendix C.

12. *Removing linearly* means that a linear classifier cannot predict the required linguistic property with above majority class accuracy.

attempts to answer this question—we provide a metric ρ that is a reliable predictor of such importance. This metric occurs naturally when we take an information-theory lens and develop a theoretical framework that ties together linguistic properties, word representations, and language modeling performance. We show that when a linguistic property is removed from word vectors, the decline in the quality of a language model depends on how strongly the removed property is interdependent with the underlying text, which is measured by ρ : *a greater ρ leads to a greater drop.*

The proposed metric has an undeniable advantage: its calculation does not require word representations themselves or a pretrained language model. All that is needed is the text and its linguistic annotation. Thanks to our Theorem 1, we can express the influence of a linguistic property on the word prediction task in terms of the coefficient ρ .

3.1 Notation

We will use plain-faced lowercase letters (x) to denote scalars and plain-faced uppercase letters (X) for random variables. Bold-faced lowercase letters (\mathbf{x}) will denote vectors—both random and non-random—in the Euclidean space \mathbb{R}^d , while bold-faced uppercase letters (\mathbf{X}) will be used for matrices.

Assuming there is a finite vocabulary \mathcal{W} , members of that vocabulary are called *tokens*. A sentence $W_{1:n}$ is a sequence of tokens $W_i \in \mathcal{W}$, this is $W_{1:n} = [W_1, W_2, \dots, W_n]$. A linguistic annotation T of a sentence $W_{1:n}$ may take different forms. For example, it may be a sequence of per-token tags $T = [T_1, T_2, \dots, T_n]$, or a parse-tree $T = (\mathcal{V}, \mathcal{E})$ with vertices $\mathcal{V} = \{W_1, \dots, W_n\}$ and edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$. We only require that T is a *deterministic* function of $W_{1:n}$.¹³

A (masked) language model is formulated as the probability distribution $q_{\theta}(W_i | \xi_i) \approx \Pr(W_i | C_i)$, where C_i is the context of W_i (see below for different types of context), and ξ_i is the vector representation of C_i . The cross-entropy loss of such a model is $\ell(W_i, \xi_i) := \mathbb{E}_{(W_i, C_i) \sim \mathcal{D}}[-\log q_{\theta}(W_i | \xi_i)]$, where \mathcal{D} is the true joint distribution of word-context pairs (W, C) .

For a random variable X , its entropy is denoted by $H[X]$. For a pair of random variables X and Y , their mutual information is denoted by $I[X; Y]$. In Appendix D we provide the necessary background on information theory, and we refer the reader to refer to it if needed.

Discreteness of representations. Depending on the LM, C_i is usually either the left context $[W_1, \dots, W_{i-1}]$ (for a causal LM) or a subsequence of the bidirectional context $[W_1, \dots, W_{i-1}, W_{i+1}, \dots, W_n]$ (for a masked LM). Although the possible set of all such contexts \mathcal{C} is infinite, it is still *countable*. Thus the set of all contextual representations $\{\xi : \xi \text{ is a vector representation of } C \mid C \in \mathcal{C}\}$ is also countable. Hence, we treat ξ as *discrete* random vector.

13. Although in reality two people can give two different annotations of the same text due to inherent ambiguity of language or different linguistic theories, we will treat T as the final—also called gold—annotation after disagreements are resolved between annotators and a common reference annotation is agreed upon.

3.2 Main Result

Our main result is the following

Theorem 1. *Let*

1. \mathbf{x}_i be a (contextualized) embedding of a token W_i in a sentence $W_{1:n}$, and denote

$$\sigma_i := I[W_i; \mathbf{x}_i] / H[W_{1:n}], \tag{1}$$

2. T be a linguistic annotation of $W_{1:n}$, and the dependence between T and $W_{1:n}$ is measured by the coefficient

$$\rho := I[T; W_{1:n}] / H[W_{1:n}], \tag{2}$$

3. $\rho > 1 - \sigma_i$,

4. $\tilde{\mathbf{x}}_i$ be a (contextualized) embedding of W_i that contains no information on T .

Then the decline in the language modeling quality when using $\tilde{\mathbf{x}}_i$ instead of \mathbf{x}_i is approximately supralinear in ρ :

$$\ell(W_i, \tilde{\mathbf{x}}_i) - \ell(W_i, \mathbf{x}_i) \gtrsim H[W_{1:n}] \cdot \rho + c \tag{3}$$

for $\rho > \rho_0$, with constants $\rho_0 > 0$, and c depending on $H[W_{1:n}]$ and $I[W_i; \mathbf{x}_i]$.

Proof. The proof is given in Appendix E. Here we provide a less formal argument. Using visualization tricks as in Olah (2015) we can illustrate the essence of the proof by Figure 6. First of all, look at the entropy $H[X]$ as the amount of information in the variable X .




Figure 6: Illustration of Theorem 1.

Imagine amounts of information as bars. These bars overlap if there is shared information between the respective variables.

The annotation T and the embedding vector $\tilde{\mathbf{x}}_i$ are derived from the underlying text $W_{1:n}$, thus $W_{1:n}$ contains more information than T or $\tilde{\mathbf{x}}_i$, hence $H[W_{1:n}]$ fully covers $H[T]$ and $H[\tilde{\mathbf{x}}_i]$. Since the mutual information $I[W_i; \tilde{\mathbf{x}}_i]$ cannot exceed the information in $\tilde{\mathbf{x}}_i$, we can write

$$I[W_i; \tilde{\mathbf{x}}_i] \leq H[\tilde{\mathbf{x}}_i]. \tag{4}$$

Now recall the Eq. (2): it simply means that ρ is the fraction of information that is left in T after it was derived from $W_{1:n}$. This immediately implies that the information which is *not* in T (but was in $W_{1:n}$ initially) is equal to $(1 - \rho) \cdot H[W_{1:n}]$.

Since the embedding $\tilde{\mathbf{x}}_i$ contains no information about the annotation T , $H[\tilde{\mathbf{x}}_i]$ and $H[T]$ do not overlap. But this means that

$$H[\tilde{\mathbf{x}}_i] \leq (1 - \rho) \cdot H[W_{1:n}], \quad (5)$$

because $\tilde{\mathbf{x}}_i$ is in the category of *no information about T*.

Combining (1), (4), (5), and the assumption $\rho > 1 - \sigma_i$, we have

$$\begin{aligned} I[W_i; \tilde{\mathbf{x}}_i] &\leq (1 - \rho) \cdot H[W_{1:n}] < \sigma_i \cdot H[W_{1:n}] = I[W_i; \mathbf{x}_i], \\ \Rightarrow I[W_i; \mathbf{x}_i] - I[W_i; \tilde{\mathbf{x}}_i] &> (\rho + \sigma_i - 1) \cdot H[W_{1:n}]. \end{aligned} \quad (6)$$

The inequality (6) is almost the required (3)—it remains to show that the change in mutual information can be approximated by the change in LM loss; this is done in Lemma E.2. \square

Role of ρ and σ_i . Equation (2) quantifies the dependence between T and $W_{1:n}$, and it simply means that the annotation T carries $100 \cdot \rho\%$ of information contained in the underlying text $W_{1:n}$ (in the information-theoretic sense). The quantity ρ is well known as the *entropy coefficient* in the information theory literature (Press et al., 2007). It can be thought of as an analog of the correlation coefficient for measuring not only linear or monotonic dependence between numerical variables but any kind of statistical dependence between any kind of variables (numerical and non-numerical). As we see from Equation (3), the coefficient ρ plays a key role in predicting the LM degradation when the linguistic structure T is removed from the embedding \mathbf{x}_i . In Section 3.3 we give a practical way of its estimation for the case when T is a per-token annotation of $W_{1:n}$.

Similarly, Equation (1) means that both W_i and \mathbf{x}_i carry at least $100 \cdot \sigma_i\%$ of information contained in $W_{1:n}$. By Firth’s distributional hypothesis (Firth, 1957),¹⁴ we assume that σ_i significantly exceeds zero.

Range of ρ . In general, mutual information is non-negative. Mutual information of the annotation T and the underlying text $W_{1:n}$ cannot exceed information contained in either of these variables, i.e. $0 \leq I[T; W_{1:n}] \leq H[W_{1:n}]$, and therefore $\rho \in [0, 1]$.

Absence of information. When we write “ $\tilde{\mathbf{x}}_i$ contains no information on T ”, this means that the mutual information between $\tilde{\mathbf{x}}_i$ and T is zero:

$$I[T; \tilde{\mathbf{x}}_i] = 0. \quad (7)$$

In the language of Pimentel et al. (2020a), Equation (7) assumes that all probes—even the best—perform poorly in extracting T from $\tilde{\mathbf{x}}_i$. This essentially means that the information on T has been filtered out of $\tilde{\mathbf{x}}$. In practice, we will approximate this with the techniques of Gradient Reversal Layer (Ganin & Lempitsky, 2015) and Iterative Nullspace Projection (Ravfogel et al., 2020).

14. “you shall know a word by the company it keeps”

3.3 Experiments

In this section, we empirically verify the prediction of our theory—the stronger the dependence of a linguistic property with the underlying text the greater the decline in the performance of a language model that does not have access to such property.

Here we will focus on only one contextualized embedding model, BERT because it is the most mainstream model. Along with this, we will keep the SGNS for consideration as a model of static embeddings.¹⁵

3.3.1 REMOVAL TECHNIQUES

We consider two ways of removing linguistic information from the embeddings: Gradient Reversal Layer (GRL) and Iterative Nullspace Projection (INLP).

GRL (Ganin & Lempitsky, 2015) is a method of adversarial training that, in our case, can be used to remove linguistic information T from word embeddings ξ . For this, the pretraining model is formulated as $\widehat{W} = \omega(\xi)$, and an auxiliary model is $\widehat{T} = \tau(\xi)$. The training procedure optimizes

$$\min_{\omega, \tau, \xi} [\ell(\omega(\xi), W) + \ell(\tau(\gamma_\lambda(\xi)), T)]$$

where $\ell(\cdot, \cdot)$ is the loss function, and γ_λ is a layer inserted between ξ and τ which acts as the identity during the forward pass, while it scales the gradients passed through it by $-\lambda$ during backpropagation. In theory, the resulting embeddings $\tilde{\mathbf{x}}$ are maximally informative for the pretraining task while at the same time minimally informative for the auxiliary task. However, in practice, GRL not always succeeds to fully remove the auxiliary information from the embeddings as was shown by Elazar and Goldberg (2018) (and confirmed by our experiments in Appendix F).

INLP (Ravfogel et al., 2020) is a method of post-hoc removal of some property T from the pretrained embeddings \mathbf{x} . INLP neutralizes the ability to linearly predict T from \mathbf{x} (here T is a single tag, and \mathbf{x} is a single vector). It does so by training a sequence of auxiliary models τ_1, \dots, τ_k that predict T from \mathbf{x} , interpreting each one as conveying information on unique directions in the latent space that correspond to T , and iteratively removing each of these directions. In the i^{th} iteration, τ_i is a *linear* model¹⁶ parameterized by a matrix \mathbf{U}_i and trained to predict T from \mathbf{x} . When the embeddings are projected onto $\text{null}(\mathbf{U}_i)$ by a projection matrix $\mathbf{P}_{\text{null}(\mathbf{U}_i)}$, we have

$$\mathbf{U}_i \mathbf{P}_{\text{null}(\mathbf{U}_i)} \mathbf{x} = \mathbf{0},$$

i.e. τ_i will be unable to predict T from $\mathbf{P}_{\text{null}(\mathbf{U}_i)} \mathbf{x}$. Figure 7 illustrates the method for the case when the property T has only two types of tags and $\mathbf{x} \in \mathbb{R}^2$. The number of iterations k is taken such that no linear classifier achieves above-majority accuracy when predicting T from $\tilde{\mathbf{x}} = \mathbf{P}_{\text{null}(\mathbf{U}_k)} \mathbf{P}_{\text{null}(\mathbf{U}_{k-1})} \dots \mathbf{P}_{\text{null}(\mathbf{U}_1)} \mathbf{x}$. Before measuring the LM loss we allow the

15. Recall that SGNS stands for skip-gram with negative sampling. SKIP-GRAM (Mikolov et al., 2013a) is a masked language model with all tokens but one in a sequence being masked. SGNS approximates its cross-entropy loss by negative sampling procedure (Bengio & Senecal, 2003).

16. Ravfogel et al. (2020) use the Linear SVM (Cortes & Vapnik, 1995), and we follow their setup.

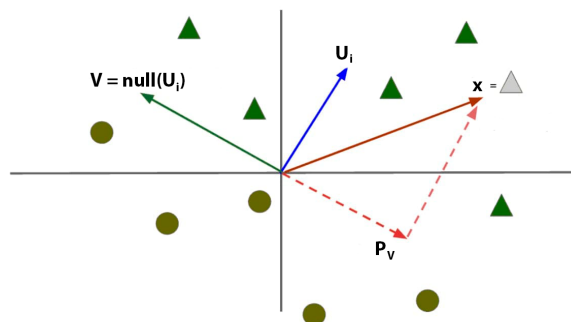


Figure 7: Nullspace projection for a 2-dimensional binary classifier. The decision boundary of \mathbf{U}_i is \mathbf{U}_i 's null-space. Source: Ravfogel et al. (2020).

model to adapt to the modified embeddings $\tilde{\mathbf{x}}$ by fine-tuning its softmax layer that predicts W from $\tilde{\mathbf{x}}$, while keeping the rest of the model (encoding layers) frozen. This fine-tuning does not bring back information on T as the softmax layer linearly separates classes.

3.3.2 TASKS

We perform experiments on two real tasks and a series of synthetic tasks. In real tasks, we cannot arbitrarily change ρ —we can only measure it for some given annotations. Although we have several real annotations that give some variation in the values of ρ , we want more control over this metric. Therefore, we come up with synthetic annotations that allow us to smoothly change ρ in a wider range and track the impact on the loss function of a language model.

To remove a property T from embeddings, INLP requires a small annotated dataset, while GRL needs the *whole* pretraining corpus to be annotated.¹⁷ Therefore for INLP we use gold annotations, while for GRL we tag the pretraining data with the help of the STANZA tagger (Qi et al., 2020).

Real tasks. We consider two real tasks with per-token annotations: part-of-speech tagging (POS) and named entity labeling (NER). The choice of per-token annotations is guided by the suggested method of estimating ρ (Section 3.3.3).

POS is the syntactic task of assigning tags such as noun, verb, adjective, etc. to individual tokens. We consider POS tagged datasets that are annotated with different tagsets:

- Universal Dependencies English Web Treebank (UD EWT, Silveira et al., 2014), which uses two annotation schemas: UPOS corresponding to 17 universal tags across languages, and FPOS which encodes additional lexical and grammatical properties.

17. One may argue that GRL can be applied to a pretrained LM in a fine-tuning mode. Our preliminary experiments show that in this case, the linear probe's accuracy remains way above the majority class accuracy. See Appendix F for details.

- English part of the OntoNotes corpus (Weischedel et al., 2013) based on the Penn Treebank annotation schema (Marcus et al., 1993).

NER is the task of predicting the category of an entity referred to by a given token, e.g. does the entity refer to a person, a location, an organization, etc. This task is taken from the English part of the OntoNotes corpus.

Note that the STANZA tagger, used to generate annotations for GRL training, relies on universal tagset from Universal Dependencies (UPOS) for POS tagging, and on OntoNotes NER annotations. Table 2 reports some statistics on the size of different tagsets.

UD UPOS	UD FPOS	ON POS	ON NER
17	50	48	66

Table 2: Size of the tagsets for Universal Dependencies and OntoNotes PoS tagging and NER datasets.

Synthetic tasks are created as follows. Let $T^{(0)}$ be an annotation of a corpus $W_{1:n}$, which has m unique tags, and let the corresponding $\rho^{(0)} := I[T^{(0)}; W_{1:n}] / H[W_{1:n}]$. We select the two least frequent tags from the tagset, and conflate them into one tag. This gives us an annotation $T^{(1)}$ which contains less information about $W_{1:n}$ than the annotation $T^{(0)}$, and thus has $\rho^{(1)} < \rho^{(0)}$. In Table 3 we give an example of such conflation for a POS-tagged sentence from the OntoNotes corpus (Weischedel et al., 2013).

$W_{1:9}$	When	a	browser	starts	to	edge	near	to	consuming
$T^{(0)}$	WRB	DT	NN	VBZ	TO	VB	RB	IN	VBG
$T^{(1)}$	X	DT	NN	VBZ	X	VB	RB	IN	VBG
$W_{10:18}$	500	MB	of	RAM	on	a	regular	basis	,
$T^{(0)}$	CD	NNS	IN	NN	IN	DT	JJ	NN	,
$T^{(1)}$	CD	NNS	IN	NN	IN	DT	JJ	NN	,
$W_{19:22}$	something	is	wrong	.					
$T^{(0)}$	NN	VBZ	JJ	.					
$T^{(1)}$	NN	VBZ	JJ	.					

Table 3: Example of conflating two least frequent tags (WRB and TO) into one tag (X).

Next, we select two least frequent tags from the annotation $T^{(1)}$ and conflate them. This will give an annotation $T^{(2)}$ with $\rho^{(2)} < \rho^{(1)}$. Iterating this process $m - 1$ times we will end up with the annotation $T^{(m-1)}$ that tags all tokens with a single (most frequent) tag. In this last iteration, the annotation has no mutual information with $W_{1:n}$, i.e. $\rho^{(m-1)} = 0$.

3.3.3 EXPERIMENTAL SETUP

We remove (pseudo-)linguistic structures (Section 3.3.2) from BERT and SGNS embeddings using the methods from Section 3.3.1,¹⁸ and measure the decline in the language modeling performance. Assuming that $\Delta\ell_{\omega,T,\mu}$ is the validation loss increase of the embedding model $\omega \in \{\text{BERT}, \text{SGNS}\}$ when information on $T \in \{\text{Synthetic}, \text{POS}, \text{NER}\}$ is removed from ω using the removal method $\mu \in \{\text{GRL}, \text{INLP}\}$, we compare $|\Delta\ell_{\omega,T,\mu}|$ against ρ defined by (2) which is the strength of interdependence between the underlying text $W_{1:n}$ and its annotation T . By Theorem 1, $\Delta\ell_{\omega,T,\mu}$ is $\Omega(\rho)$ ¹⁹ for any combination of ω , T , μ .

Estimating ρ . Recall that $\rho := I[T; W_{1:n}] / H[W_{1:n}]$ (Eq. 2) and that the annotation T is a deterministic function of the underlying text $W_{1:n}$ (Sec. 3.1). In this case, we can write²⁰

$$\rho = \frac{H[T] - \overbrace{H[T | W_{1:n}]}^0}{H[W_{1:n}]} = \frac{H[T]}{H[W_{1:n}]} \quad (8)$$

and when T is a per-token annotation of $W_{1:n}$, i.e. $T = T_{1:n}$ (which is the case for the annotations that we consider), this becomes $\rho = H[T_{1:n}] / H[W_{1:n}]$. Thus to estimate ρ , we simply need to be able to estimate the latter two entropies. This can be done by training an autoregressive sequence model, such as LSTM, on $W_{1:n}$ and on $T_{1:n}$. The loss function of such a model—the cross-entropy loss—serves as an estimate of the required entropy. Notice, that we cannot use masked LMs for this estimation as they do not give a proper factorization of the probability $p(w_1, \dots, w_n)$. Thus, we decided to choose the AWD-LSTM-MoS model of Yang et al. (2018) which is a compact and competitive LM²¹ that can be trained in a reasonable time and with moderate computing resources. In addition, we also estimated the entropies through a vanilla LSTM with tied input and output embeddings (Inan et al., 2017), and a Kneser-Ney 3-gram model²² (Ney et al., 1994) to test how strongly our method depends on the underlying sequence model.

Limitations. The suggested method of estimating ρ through autoregressive sequence models is limited to per-token annotations only. However, according to formula (8), to estimate ρ for deeper annotations, it is sufficient to be able to estimate the entropy $H[T]$ of such deeper linguistic structures T . For example, to estimate the entropy of a parse tree, one can use the cross-entropy produced by a probabilistic parser. The only limitation is the determinism of the annotation process.

Amount of information to remove. Each of the removal methods has a hyperparameter that controls how much of the linguistic information T is removed from the word vectors \mathbf{x} —the number of iterations in INLP and λ in GRL. Following Elazar et al. (2021) we keep

18. INLP is applied to the last layers of BERT and SGNS.

19. as a reminder, this means that there exist constants c_0, c_1, ρ_0 s.t. $\Delta\ell_{\omega,T,\mu} \geq c_1 \cdot \rho + c_0$ for $\rho \geq \rho_0$.

20. In general, for two random variables X and Y , $Y = f(X)$ if and only if $H[Y | X] = 0$.

21. <https://paperswithcode.com/sota/language-modelling-on-wikitext-2>

22. We used the SRILM toolkit (Stolcke, 2002). We used Witten-Bell discounting for annotations because modified Kneser-Ney discounting does not apply to them. This is a known issue with Kneser-Ney when the vocabulary size (number of unique tags) is small. See item C3 at <http://www.cs.cmu.edu/afs/cs/project/cmt-55/lti/Courses/731/homework/srilm/man/html/srilm-faq.7.html>.

iterating the INLP procedure or increasing λ in GRL until the performance of a linear probe that predicts T from the filtered embeddings $\tilde{\mathbf{x}}$ drops to the majority-class accuracy. When this happens we treat the resulting filtered embeddings $\tilde{\mathbf{x}}$ as containing no information on T .

Optimization details. For the INLP experiments we use pretrained BERT-BASE from HuggingFace (Wolf et al., 2020), and an SGNS pretrained in-house (Assylbekov, 2020). For the GRL experiments we pretrain LMs ourselves, details are in Appendix A.

3.3.4 RESULTS

Real tasks. Table 4 reports the loss drop of pretrained LM when removing linguistic information (POS or NER) from the pretrained model. Plots that illustrate how losses and probing accuracies change with INLP iterations are provided in Appendix G.

First, we compare UPOS tagsets versus FPOS tagset: intuitively FPOS should have a tighter link with underlying text, and therefore result in higher ρ and as a consequence in a higher drop in loss after the removal of this information from words representations. This is confirmed by the numbers reported in Table 4.

We also see a greater LM performance drop when the POS information is removed from the models compared to NE information removal. This is in line with Theorem 1 as POS tags depend stronger on the underlying text than NE labels as measured by ρ .

In the case of pretraining SGNS with GRL, we could not reach the majority class accuracy.²³ Reported is the loss increase for the model that had the lowest accuracy (Acc = .466 while the majority accuracy for this task is .126). If it were possible to train the SGNS model with GRL removing POS tags and reaching the majority class accuracy, its loss increase would be even higher than the one reported.

Removal method μ Annotation T	INLP			GRL		
	ON NER	UD UPOS	UD FPOS	ON POS	NER	UPOS
ρ (AWD-LSTM-MoS)	0.18	0.32	0.36	0.42	0.18	0.43
ρ (LSTM)	0.18	0.32	0.37	0.42	0.22	0.34
ρ (KN-3)	0.18	0.36	0.41	0.42	0.17	0.37
$\Delta\ell_{\text{BERT},T,\mu}$	0.13	0.54	0.70	0.87	6.16	6.16
$\Delta\ell_{\text{SGNS},T,\mu}$	1.33	1.62	1.79	2.04	0.32	>1.1 [†]

Table 4: Results on POS tagging and NER tasks. ON stands for OntoNotes, UD for Universal Dependencies. The UD EWT dataset has two types of POS annotation: coarse tags (UPOS) and fine-grained tags (FPOS). [†]GRL for SGNS on UPOS annotation did not reach the majority class accuracy (.126). KN-3 is a Kneser-Ney 3-gram model.

We note larger $\Delta\ell_{\text{BERT},T,\text{GRL}}$ compared to $\Delta\ell_{\text{BERT},T,\text{INLP}}$. We attribute it to the following: INLP works with already pretrained *optimal* ω and we fine-tune its softmax layer

23. This is because GRL does not always fully remove all the auxiliary task information from the main model (Elazar & Goldberg, 2018).

for modified embeddings $\tilde{\mathbf{x}}$, while GRL with a large λ (that is needed to bring the probing accuracy down to a majority class) moves ω 's weights into a bad local minimum when pre-training from scratch. Moreover, the pretrained BERT from HuggingFace and ROBERTA pretrained by us are pretrained over different corpora and are using different vocabularies. Thus increases in the cross-entropies caused by INLP and by GRL are not directly comparable.

We attempted to conduct experiments on GRL in such a way that their results were comparable to the results of the INLP—we tried to fine-tune a pretrained RoBERTa with GRL. However, our attempt was not successful as we did not manage to reach majority class accuracy (more details in Appendix F). This reinforces Elazar and Goldberg's (2018) argument that GRL has more difficulties removing fully the auxiliary information from the word representations.

Finally, we see that although ρ indeed depends on the underlying sequence model that is used to estimate the entropies $H[T_{1:n}]$ and $H[W_{1:n}]$, all models—AWD-LSTM-MoS, LSTM, and KN-3—preserve the relative order for the annotations that we consider. E.g., all models indicate that OntoNotes NER annotation is the least interdependent with the underlying text, while OntoNotes POS annotation is the most interdependent one. In addition, it turns out that for a quick estimate of ρ , one can use the KN-3 model, which on a modern laptop calculates the entropy of texts of 100 million tokens in a few minutes, in contrast to the LSTM, which takes several hours on a modern GPU.

Synthetic tasks. To obtain synthetic data, we apply the procedure described in Sect. 3.3.2 to the OntoNotes POS annotation as it has the highest ρ in Table 4 and thus allows us to vary the metric in a wider range.

The results of evaluation on the synthetic tasks through the INLP²⁴ are provided in Figure 8. They validate the predictions of our theory—for the annotations with greater ρ there is a bigger drop in the LM performance (i.e. increase in the LM loss) when the information on such annotations is removed from the embeddings. We notice that $\Delta\ell$ is piecewise-linear in ρ with the slope changing at $\rho \approx 0.4$. We attribute this change to the following: for $\rho < 0.4$, the majority class (i.e. the most frequent tag) is the tag that encapsulates several conflated tags (see Subsection 3.3.2 for details), while for $\rho > 0.4$, the majority is NN tag. This switch causes a significant drop in the majority class accuracy which in turn causes a significant increase in the number of INLP iterations to reach that accuracy, and hence an increase in the amount of information being removed which implies greater degradation of the LM performance.

4. Related Work

Theoretical analysis. Since the success of early word embedding algorithms like SGNS and GLOVE, there were attempts to analyze theoretically the connection between their pretraining objectives and performance on downstream tasks such as word similarity and word analogy tasks. An incomplete list of such attempts includes those of Levy and Goldberg (2014), Arora et al. (2016), Hashimoto et al. (2016), Gittens et al. (2017), Tian et al. (2017),

24. We did not perform GRL experiments on synthetic tasks as they are computationally expensive—finding λ for each annotation requires multiple pretraining runs.

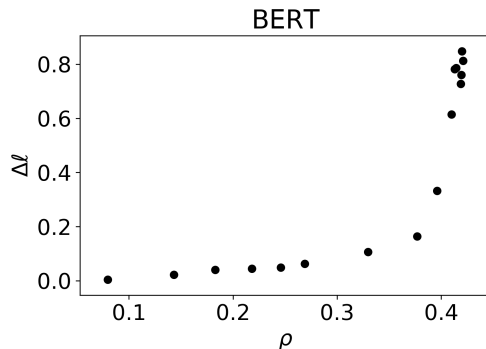


Figure 8: Synthetic task results for INLP. $\Delta \ell$ is the increase in cross-entropy loss when pseudo-linguistic information is removed from the BERT’s last layer with the INLP procedure. ρ is estimated with the help of AWD-LSTM-MoS model (Yang et al., 2018) as described in Section 3.3.3.

Ethayarajh et al. (2019), Allen and Hospedales (2019). Most of these works represent pretraining as a low-rank approximation of some co-occurrence matrix—such as PMI—and then use an empirical fact that the set of columns (or rows) of such a matrix is already a good solution to the analogy and similarity tasks. Recently, we have seen a growing number of works devoted to the theoretical analysis of contextualized embeddings. Kong et al. (2020) showed that modern embedding models, as well as the old warrior SGNS, maximize an objective function that is a lower bound on the mutual information between different parts of the text. Lee et al. (2020) formalized how solving certain pretraining tasks allows learning representations that provably decrease the sample complexity of downstream supervised tasks. Of particular interest is a recent paper by Saunshi et al. (2021) that relates a pretraining performance of an autoregressive LM with a downstream performance for downstream tasks that *can* be reformulated as next word prediction tasks. The authors showed that for such tasks, if the pretraining objective is ϵ -optimal,²⁵ then the downstream objective of a linear classifier is $\mathcal{O}(\sqrt{\epsilon})$ -optimal. In the second part of our work (Section 3) we prove a similar statement, but the difference is that we study how the removal of linguistic information affects the pretraining objective and our approach is not limited to downstream tasks that can be reformulated as next word prediction.

Probing. Early work on probing tried to analyze LSTM language models (Linzen et al., 2016; Shi et al., 2016; Adi et al., 2017; Conneau et al., 2018; Gulordava et al., 2018; Kuncoro et al., 2018; Tenney et al., 2019b). Moreover, word similarity (Finkelstein et al., 2002) and word analogy (Mikolov et al., 2013d) tasks can be regarded as non-parametric probes of static embeddings such as SGNS (Mikolov et al., 2013b) and GLOVE (Pennington et al., 2014). Recently the probing approach has been used mainly for the analysis of contextualized word embeddings. Hewitt and Manning (2019) for example showed that entire parse trees can be linearly extracted from ELMo’s (Peters et al., 2018) and BERT’s (Devlin et al.,

25. Saunshi et al. (2021) say that the pre-training loss ℓ is ϵ -optimal if $\ell - \ell^* \leq \epsilon$, where ℓ^* is the minimum achievable loss.

2019) hidden layers. Tenney et al. (2019b) probed contextualized embeddings for various linguistic phenomena and showed that, in general, contextualized embeddings improve over their non-contextualized counterparts largely on syntactic tasks (e.g., constituent labeling) in comparison to semantic tasks (e.g., coreference). The probing methodology has also shown that BERT learns some reflections of semantics (Reif et al., 2019) and factual knowledge (Petroni et al., 2019b) into the linguistic form which are useful in applications such as word sense disambiguation and question answering respectively. Zhang et al. (2020) analyzed how the quality of representations in a pretrained model evolves with the amount of pretraining data. They performed extensive probing experiments on various NLU tasks and found that pretraining with 10M sentences was already able to solve most of the syntactic tasks, while it required 1B training sentences to be able to solve tasks requiring semantic knowledge (such as Named Entity Labeling, Semantic Role Labeling, and some others as defined by Tenney et al. (2019b)).

Closest to our second part (Section 3), Elazar et al. (2021) propose to look at probing from a different angle, proposing *amnesic probing* which is defined as the drop in performance of a pretrained LM after the relevant linguistic information is removed from one of its layers. The notion of amnesic probing fully relies on the assumption that the amount of the linguistic information contained in the pretrained vectors should correlate with the drop in LM performance after this information is removed. In this work (Section 3) we theoretically prove this assumption. While Elazar et al. (2021) measured LM performance as word prediction accuracy, we focus on the native LM cross-entropy loss. In addition, we answer one of the questions raised by the authors on how to measure the influence of different linguistic properties on the word prediction task—we provide an easy-to-estimate metric that does exactly this.

Criticism of probing. The probing approach has been criticized from different angles. Our attempt to systematize this line of work is given in Figure 9. The seminal paper of Hewitt and Liang (2019) raises the issue of separation between *extracting* linguistic structures from contextualized embeddings and *learning* such structures by the probes themselves. This dichotomy was challenged by Pimentel et al. (2020b), Maudslay et al. (2020), but later validated by Zhu and Rudzicz (2020) using an information-theoretic view on probing. Meanwhile, methods were proposed that take into account not only the probing performance but also the ease of extracting linguistic information (Voita & Titov, 2020) or the complexity of the probing model (Pimentel et al., 2020a). At the same time, Wu et al. (2020) and Michael et al. (2020) suggested avoiding learnability issues by non-parametric probing²⁶ and weak supervision respectively. The remainder of the criticism is directed at the limitations of probing such as insufficient reliability for low-resourced languages (Eger et al., 2020), lack of evidence that probes indeed extract linguistic structures but do not learn from the linear context only (Kunz & Kuhlmann, 2020), lack of correlation with fine-tuning scores (Tamkin et al., 2020) and with pretraining scores (Ravichander et al., 2020; Elazar et al., 2021). The first part of our work (Section 2) partly falls into this latter group, as we did not find

26. Parametric probes transform embeddings to linguistic structures using *parameterized* operations on vectors (such as feedforward layers). Non-parametric probes transform embeddings using *non-parameterized* operations on vectors (such as vector addition/subtraction, inner product, Euclidean distance, etc.). The approach of Wu et al. (2020) builds a so-called impact matrix and then feeds it into a graph-based algorithm to induce a dependency tree, all done without learning any parameters.

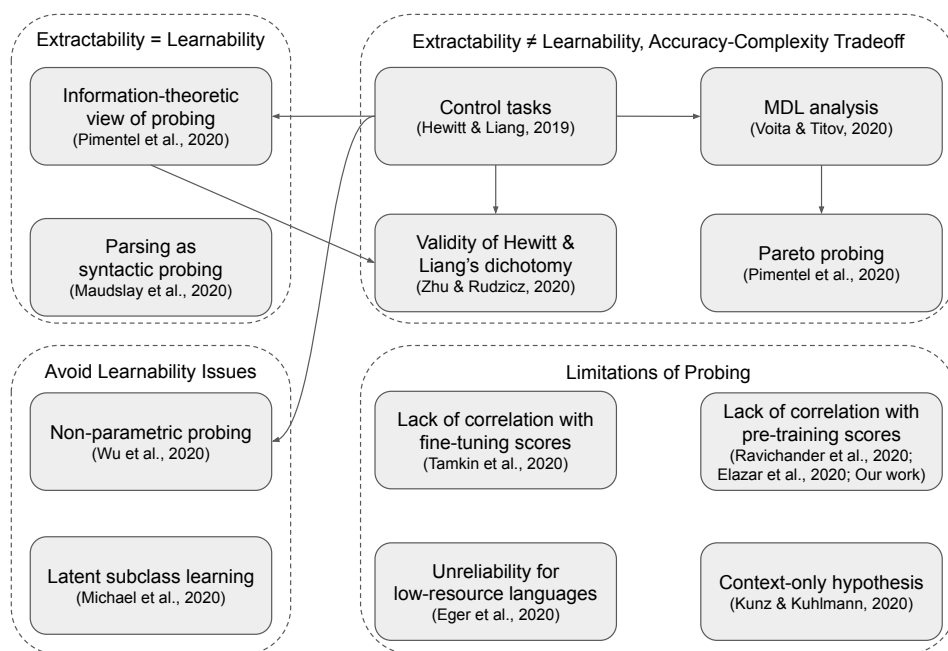


Figure 9: Criticism and improvement of the probing methodology. An arrow $A \rightarrow B$ means that B criticizes and/or improves A .

any evidence for a correlation between probing scores and pretraining objectives for better performing CoVe (McCann et al., 2017).

Pruning language models. A recent work by Gordon et al. (2020) compressed BERT using conventional pruning and showed a linear correlation between pretraining loss and downstream task accuracy. Chen et al. (2020) pruned pretrained BERT with LTH and fine-tuned it to downstream tasks, while Prasanna et al. (2020) pruned fine-tuned BERT with LTH and then re-fine-tuned it. Sanh et al. (2020) showed that the weights needed for specific tasks are a small subset of the weights needed for masked language modeling, but they prune *during fine-tuning* which is beyond the scope of our work. Zhao et al. (2020) propose to learn the masks of the pretrained LM as an alternative to finetuning on downstream tasks and shows that it is possible to find a subnetwork of large pretrained model which can reach the performance on the downstream tasks comparable to finetuning on this task. Generally speaking, the findings of the above-mentioned papers are aligned with our findings that the performance of pruned models on downstream tasks is correlated with the pretraining loss. The one difference from Section 2 of our work is that most of the previous work looks at the performance of *fine-tuned* pruned models. In our work, we *probe* pruned models, i.e. the remaining weights of language models are *not* adjusted to the downstream probing task. It is not obvious whether the conclusions from the former should carry over to the latter.

5. Conclusion

In this work, we tried to better understand the phenomenon of the *rediscovery hypothesis* in pretrained language models. Our main contribution is two-fold: we demonstrate that the *rediscovery hypothesis*

- holds across various pretrained LMs (SGNS, CoVE, RoBERTa) even when a significant amount of weights get pruned (in English);
- can be formally defined within an information-theoretic framework and proved (assuming that the linguistic annotation is a deterministic function of the underlying text and that the annotation is sufficiently interdependent with the text).

First (Section 2), we performed probing of different pruned²⁷ instances of the original models. If models are overparametrized, then it could be that the pruned model only keeps the connections that are important for the pretraining task, but not for auxiliary tasks like probing. Our experiments show that there is a correlation between the pretrained model’s cross-entropy loss and probing performance on various linguistic tasks. We believe that such correlation can be interpreted as strong evidence in favor of the *rediscovery hypothesis*.

While performing the above-mentioned analysis, several additional interesting observations have been made.

- Break-down frequency analysis (Section 2.5) reveals that at moderate pruning rates the POS tag performance drop is quite similar across all the token frequency groups. It indicates that to preserve LM performance, the pruned models were forced to distribute their “linguistic knowledge” across all the token groups, and thus reinforces the *rediscovery hypothesis*.

We noted, however, that part of the probing performance could be due to exploitation of *spurious memorization* in pretraining LM. A question on whether such *spurious memorization* should be considered as *linguistic knowledge* learnt by pretrained LM is an ongoing debate in the community (Bender & Koller, 2020; Sahlgren & Carlsson, 2021). We believe that further investigation in this direction could bring more insights into the pretrained model’s behavior.

- We have analyzed how the amount of *linguistic knowledge* evolves during the pre-training process. Our analysis of intermediate pretraining checkpoints suggests that models follow *NLP pipeline discovery during* the pretraining process: it reaches good performance on syntactic tasks at the beginning of the training and keeps improving on semantic tasks further in training.

Section 3 proposes an information-theoretic formulation of the rediscovery hypothesis and a measure of the dependence between the text and linguistic annotations. In particular, Theorem 1 states that neural language models²⁸ *do need to discover linguistic information* to perform at their best: when the language models are denied access to linguistic information,

27. Pruning was done through iterative magnitude pruning (aka Lottery Ticket Hypothesis) or 1-shot magnitude pruning

28. This is, language models which rely on intermediate representations of tokens.

their performance is not optimal. The *metric* ρ , which naturally appears in the development of our information-theoretic framework proposes a measure of the relationship between a given linguistic property and the underlying text. We show that it is also a reliable predictor of how suboptimal a language model will be if it is denied access to a given linguistic property. Our proposed methods of estimating ρ can be used in further studies to *quantify* the dependence between text and its per-token annotation.

A number of previous studies (Tenney et al., 2019b; Conneau et al., 2018; Hewitt & Manning, 2019; Zhang et al., 2020) have shown some evidence in favour of *rediscovery hypothesis*, while many others (Voita & Titov, 2020; Pimentel et al., 2020b; Hewitt & Liang, 2019; Elazar et al., 2021) questioned the methodology of those studies. However, none of the above-mentioned works questioned the rediscovery hypothesis *per se* nor demonstrated any evidence for rejecting it. We believe that our work brings an additional understanding of this phenomenon. In addition, it is the first attempt to formalize the rediscovery hypothesis and provide a theoretical frame for it.

Acknowledgements

The authors would like to thank the anonymous JAIR reviewers of this paper, the reviewers of previous versions, as well as Marc Dymetman, Arianna Bisazza, Germán Kruszewski, Laurent Besacier, Shauli Ravfogel, and Yanai Elazar, for their detailed and constructive feedback. The work of Maxat Tezekbayev, Nuradil Kozhakhmet, Madina Babazhanova, and Zhenisbek Assylbekov is supported by the Nazarbayev University Collaborative Research Program 091019CRP2109.

Appendix A. Optimization

CoVe

Two-layer BiLSTM encoder with size 600. For embedding matrix, we use GLOVE embeddings with 840B tokens and fix embedding matrix while training as was done in the original implementation (McCann et al., 2017). Thus we do not prune the embedding matrix in LTH experiments. For training, we used IWSLT 16 dataset. Scripts for preprocessing the data and training the model were adapted to a newer version of OPENNMT-PY (Klein et al., 2017). We use the default hyperparameters of OPENNMT-PY: batch size 64, maximum training steps 100,000. We use the following early stopping criterion: we stop training the model if there is no improvement in 3 consecutive validation perplexity scores (validation is performed every 2,000 steps).

RoBERTa

Optimization is performed on the full Wikitext-103 data. We used the default `roberta.base` tokenizer (distributed along with `roberta.base` model) (Liu et al., 2019). The following training settings were used for training on 2 GPUs for 100 epochs:

- architecture: embedding size 512, 6 layers, 4 heads, hidden size 1024
- batch size: 4096 tokens (with gradient accumulation for 4 steps)




Figure 10: Zooming in on better performing CoVE models. Indication of the region to be zoomed in (left) and the zoomed in region (right).

- dropout 0.3

Pruning rates used for LTH experiments are chosen as follows:

- step 0.025 for pruning up to 20% of model parameters
- step 0.05 between 20% until 70% of model parameters

SGNS

Embedding size 200, 15 epochs to train, 5 negative samples per training example, batch size 1024, window size 5. For training we use the `text8` dataset (Mahoney, 2011), which is the first 100MB of the English Wikipedia dump on Mar. 3, 2006. For validation we use the next 10MB of the same dump. We were ignoring words that appeared less than 5 times in `text8`, and subsampling the rest as follows: each word w in training set is discarded with probability $p = \sqrt{t/f(w)} + t/f(w)$, where $t = 10^{-4}$, and $f(w)$ is the frequency of word w . We use ADAM optimizer with its default settings.

Appendix B. Lack of Correlation for Better Performing Models

In Section 2 we noticed that when restricting the loss values of CoVE to the better performing end (zoomed in regions in Figure 10), drop in probing scores (compared to the scores of the baseline unpruned model) become indistinguishable. Here we treat this claim more rigorously. Assuming that ℓ_i^{CoVE} is the validation loss of CoVE at iteration i , and $\Delta s_i^{\text{CoVE}, T}$ is the corresponding drop in the score on the probing task $T \in \{\text{NE}, \text{POS}, \text{Constituents}, \text{Structural}\}$, we obtain pairs $(\ell_i^{\text{CoVE}}, \Delta s_i^{\text{CoVE}, T})$ and setup a simple linear regression for the zoomed in regions in Figure 10:

$$\Delta s_i^{\text{CoVE}, T} = \alpha + \beta \cdot \ell_i^{\text{CoVE}} + \epsilon_i.$$

Below are p -values when testing $H_0 : \beta = 0$ with a two-sided Student t -test. In all cases β is *not* significantly different from 0. Hence, the correlation between model performance and its linguistic knowledge may become very weak for better solutions to the pretraining objectives.

Task	NE	POS	Constituents	Structural
p-value	0.725	0.306	0.184	0.160

Table 5: Testing $\beta = 0$, reported are p -values.

We argue that this phenomenon is related to the amount of information remaining in an LM after pruning, and how much of this information is enough for the probe. It turns out that in the best CoVE (with the lowest pretraining loss) such information is redundant for the probe, and in the not-so-good (but still decent) CoVE there is just enough of information for the probe to show its best performance.

MDL Analysis

Recently, Voita and Titov (2020) have addressed a related issue by proposing a minimum description length (MDL) analysis of probes, that takes into account the *difficulty* of extracting the information. Here, we apply their analysis to CoVE (that is faster to experiment with than ROBERTA), on top of the edge probe for constituents.²⁹ The resulting online code-lengths vs cross-entropy values are provided in Fig. 11. We observe an even *wider* spectrum

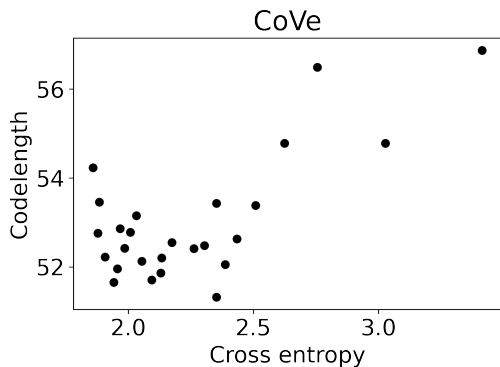


Figure 11: MDL probing of CoVE for constituents (online code). Unlike F1 score, the code length takes into account not only the performance but also the amount of efforts to achieve it.

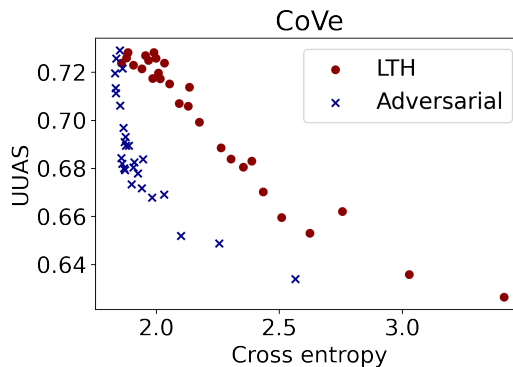


Figure 12: Structural probing of adversarially trained CoVE.

of models (with cross-entropies 1.9–2.5) that have indistinguishable code lengths. Nevertheless, the overall pattern does not contradict the rediscovery hypothesis—better models have better probing scores (lower code lengths).

²⁹. MDL requires the probe’s response to be categorical and thus cannot be applied directly to the structural probe that predicts tree-distance between the vertices of a parse tree.

Appendix C. Adversarial Training

In Sect. 2.3 we have analyzed linguistic knowledge discovered through standard language modeling pretraining. Here we propose a more radical approach, consisting in explicitly adding adversarial noise to the pretraining objective of COVE that enforces *random* angles between contextualized representations of nearby words in the source sentence of the translation task. Different to what is done in Sect. 3, this is done on COVE. For this, we exploit an empirical fact that the cosine similarity between word vectors tends to be a good predictor of their semantic similarity when trained non-adversarially. We vary the amount of noise (through a hyperparameter), train the models, and probe them with the structural probe.

C.1 Training Details

Let $F = [f_1, f_2, \dots, f_n]$ be a sentence in the source language, and $E = [e_1, e_2, \dots, e_m]$ be a corresponding sentence in the target language. Parallel corpus is the set of sentence pairs $\mathcal{D} = \{(F, E)\}$. Neural machine translation (NMT) model parameterized by θ solves

$$\min_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(F,E) \in \mathcal{D}} \underbrace{-\log p_{\theta}(E | F)}_{\ell^{\text{NMT}}(F,E)} \quad (9)$$

In case of COVE, the NMT model is an LSTM-based encoder-decoder architecture with attention (Bahdanau et al., 2015), and let $\mathbf{h}_{1:n} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n]$ be the top-level activations of the encoder that correspond to the sentence F .

Recall that for a pair of words f_i and f_j , the cosine similarity between their vector representations \mathbf{h}_i and \mathbf{h}_j strongly correlates with the human similarity rating $r(f_i, f_j)$ (Mikolov et al., 2013b). Strong correlation means that there is an approximately linear relationship between $r(f_i, f_j)$ and $\cos(\mathbf{h}_i, \mathbf{h}_j)$. In order to “break” this relationship we fit $\cos(\mathbf{h}_i, \mathbf{h}_j)$ to a *random* $r \sim \mathcal{U}(0, 1)$ for nearby source words while training the NMT model (which COVE model is based on). For a pair (F, E) an individual objective becomes

$$\ell(F, E) = \ell^{\text{NMT}}(F, E) + \alpha \cdot \ell^{\text{Adv}}(F),$$

where

$$\ell^{\text{Adv}}(F) = \frac{1}{n} \sum_{i=1+k}^{n-k} \sum_{\substack{i-k \leq j \leq i+k \\ j \neq i}} \mathbb{E}_{r_{ij} \sim \mathcal{U}(0,1)} [w \cdot \cos(\mathbf{h}_i, \mathbf{h}_j) + b - r_{ij}]^2,$$

w and b are regression coefficients, k is the size of a sliding window, in our experiments we use $k = 5$. Hyperparameter α regulates the impact of the adversarial loss.

C.2 Results

The probing results are provided in Fig. 12, where cross entropies are COVE’s validation loss values without the adversarial additive. Although the probing score decreases with an increase of cross-entropy, it does so significantly faster than in the case of LTH—we can have a near-to-best LM with a mediocre probing score. This would suggest that it is possible to find the state of the network that would encode less linguistic information given the same

pretraining loss. This is however not the path naturally chosen by the pretraining models, therefore probably more difficult to reach. Note, that several artifacts of CoVE training could also explain these results. First, it can be because the pretrained embedding layer is not updated in these experiments, therefore keeping a good (static) representation of tokens. Moreover, recall that CoVE is an encoder-decoder architecture, and reported cross-entropy loss is computed based on the decoder representations. The adversarial loss however has been defined on the encoder representations. Thus, even if the encoder representations got corrupted, the decoder still has a chance to recover and obtain reasonable cross-entropy loss. More extensive adversarial experiments might be necessary to be able to make any conclusion from this experiment.

Appendix D. Information Theory Background

For a random variable X with the distribution function $p(x)$, its *entropy* is defined as

$$H[X] := \mathbb{E}_X[-\log p(X)].$$

The quantity inside the expectation, $-\log p(x)$, is usually referred to as *information content* or *surprisal*, and can be interpreted as quantifying the level of “surprise” of a particular outcome x . As $p(x) \rightarrow 0$, the surprise of observing x approaches $+\infty$, and conversely as $p(x) \rightarrow 1$, the surprise of observing x approaches 0. Then the entropy $H[X]$ can be interpreted as the average level of “information” or “surprise” inherent in the X ’s possible outcomes. We will use the following

Property D.1. *If $Y = f(X)$, then*

$$H[Y] \leq H[X], \tag{10}$$

i.e. the entropy of a variable can only decrease when the latter is passed through a (deterministic) function.

Now let X and Y be a pair of random variables with the joint distribution function $p(x, y)$. The *conditional entropy* of Y given X is defined as

$$H[Y | X] := \mathbb{E}_{X,Y}[-\log p(Y | X)],$$

and can be interpreted as the amount of information needed to describe the outcome of Y given that the outcome of X is known. The following property is important for us.

Property D.2. $H[Y | X] = 0 \Leftrightarrow Y = f(X)$.

The *mutual information* of X and Y is defined as

$$I[X; Y] := \mathbb{E}_{X,Y} \left[\log \frac{p(X, Y)}{p(X) \cdot p(Y)} \right]$$

and is a measure of mutual dependence between X and Y . We will need the following properties of the mutual information.

Property D.3. 1. $I[X; Y] = H[X] - H[X | Y] = H[Y] - H[Y | X]$.

2. $I[X; Y] = 0 \Leftrightarrow X$ and Y are independent.
3. For a function f , $I[X; Y] \geq I[X; f(Y)]$.
4. $H[X, Y] = H[X] + H[Y] - I[X; Y]$,

Property D.3.3 is known as data processing inequality, and it means that post-processing cannot increase information. By *post-processing* we mean a transformation $f(Y)$ of a random variable Y , independent of other random variables. In Property D.3.4, $H[X, Y]$ is the *joint entropy* of X and Y which is defined as

$$H[X, Y] := \mathbb{E}_{X, Y}[-\log p(X, Y)]$$

and is a measure of “information” associated with the outcomes of the tuple (X, Y) .

Appendix E. Proof of Theorem 1

The proof is split into two steps: first, in Lemma E.1 we show that the decrease in mutual information $\Delta I := I[W_i; \mathbf{x}_i] - I[W_i; \tilde{\mathbf{x}}_i]$ is $\Omega(\rho)$, and then in Lemma E.2 we approximate ΔI by the increase in cross-entropy loss $\ell(W_i, \tilde{\mathbf{x}}_i) - \ell(W_i, \mathbf{x}_i)$.

Lemma E.1. *Let T be an annotation of a sentence $W_{1:n} = [W_1, \dots, W_i, \dots, W_n]$ and let $\mathbf{x}_i, \tilde{\mathbf{x}}_i$ be (contextualized) word vectors corresponding to a token W_i such that*

$$I[T; \tilde{\mathbf{x}}_i] = 0. \quad (11)$$

Denote

$$\rho := I[T; W_{1:n}] / H[W_{1:n}], \quad \rho \in [0, 1], \quad (12)$$

$$\sigma_i := I[W_i; \mathbf{x}_i] / H[W_{1:n}], \quad \sigma_i \in [0, 1]. \quad (13)$$

If $\rho > 1 - \sigma_i$, then

$$I[W_i; \tilde{\mathbf{x}}_i] < I[W_i; \mathbf{x}_i]. \quad (14)$$

and the difference $\Delta I := I[W_i; \mathbf{x}_i] - I[W_i; \tilde{\mathbf{x}}_i]$ is lowerbounded as

$$\Delta I \geq (\rho + \sigma_i - 1) \cdot H[W_{1:n}]. \quad (15)$$

Proof. From (11) and by Property D.3.4, we have

$$\begin{aligned} H[T, \tilde{\mathbf{x}}_i] &= H[T] + H[\tilde{\mathbf{x}}_i] - \underbrace{I[T; \tilde{\mathbf{x}}_i]}_0 \\ &= H[T] + H[\tilde{\mathbf{x}}_i]. \end{aligned} \quad (16)$$

On the other hand, both the annotation T and the word vectors $\tilde{\mathbf{x}}_i$ are obtained from the underlying sentence $W_{1:n}$, and therefore $T = T(W_{1:n})$ and $\tilde{\mathbf{x}}_i = \tilde{\mathbf{x}}_i(W_{1:n})$, i.e. the tuple $(T, \tilde{\mathbf{x}}_i)$ is a function of $W_{1:n}$, and by Property D.1,

$$H[T, \tilde{\mathbf{x}}_i] \leq H[W_{1:n}]. \quad (17)$$

From (16) and (17), we get

$$\mathbb{H}[T] + \mathbb{H}[\tilde{\mathbf{x}}_i] \leq \mathbb{H}[W_{1:n}]. \quad (18)$$

Since $T = T(W_{1:n})$ and $\tilde{\mathbf{x}}_i = \tilde{\mathbf{x}}_i(W_{1:n})$, by Property D.2 we have $\mathbb{H}[T | W_{1:n}] = 0$ and $\mathbb{H}[\tilde{\mathbf{x}}_i | W_{1:n}] = 0$, and therefore

$$\mathbb{I}[T; W_{1:n}] = \mathbb{H}[T] - \mathbb{H}[T | W_{1:n}] = \mathbb{H}[T]$$

Plugging this into (18), rearranging the terms, and taking into account (12), we get

$$\mathbb{H}[\tilde{\mathbf{x}}_i] \leq \mathbb{H}[W_{1:n}] - \underbrace{\mathbb{I}[T; W_{1:n}]}_{\rho \cdot \mathbb{H}[W_{1:n}]} = (1 - \rho) \cdot \mathbb{H}[W_{1:n}]. \quad (19)$$

Also, by Property D.3.1,

$$\mathbb{I}[W_i; \tilde{\mathbf{x}}_i] = \mathbb{H}[\tilde{\mathbf{x}}_i] - \underbrace{\mathbb{H}[\tilde{\mathbf{x}}_i | W_i]}_{\geq 0} \leq \mathbb{H}[\tilde{\mathbf{x}}_i] \quad (20)$$

From (13), (19), (20), and the assumption $\rho > 1 - \sigma_i$, we have

$$\mathbb{I}[W_i; \tilde{\mathbf{x}}_i] \leq \mathbb{H}[\tilde{\mathbf{x}}_i] \leq (1 - \rho) \cdot \mathbb{H}[W_{1:n}] < \sigma_i \cdot \mathbb{H}[W_{1:n}] = \mathbb{I}[W_i; \mathbf{x}_i],$$

which implies (14) and (15). \square

Remark 1. Equation (12) quantifies the dependence between T and $W_{1:n}$, and it means that T carries $100 \cdot \rho\%$ of information contained in $W_{1:n}$. Similarly, Equation (13) means that both W_i and \mathbf{x}_i carry at least $100 \cdot \sigma_i\%$ of information contained in $W_{1:n}$. By Firth’s distributional hypothesis (Firth, 1957),³⁰ σ_i significantly exceeds zero.

Remark 2. The mutual information $\mathbb{I}[W_i; \boldsymbol{\xi}_i]$ can be interpreted as the performance of the best language model that tries to predict the token W_i given its contextual representation $\boldsymbol{\xi}_i$.³¹ Therefore, inequalities (14) and (15) mean that removing a linguistic structure from word vectors *does* harm the performance of a language model based on such word vectors, and the more the linguistic structure is interdependent with the underlying text (that is being predicted by the language model) the bigger is the harm.

We treat $\mathbb{I}[W; \boldsymbol{\xi}]$ as the performance of a language model. However, in practice, this performance is measured by the validation loss ℓ of such a model, which is usually the cross-entropy loss. Nevertheless, the change in mutual information *can* be estimated by the change in the LM objective, as we show below.

Lemma E.2. *Let \mathbf{x}_i and $\tilde{\mathbf{x}}_i$ be (contextualized) embeddings of a token W_i in a sentence $W_{1:n} = [W_1, \dots, W_i, \dots, W_n]$. Let $\ell(W_i, \boldsymbol{\xi}_i)$ be the cross-entropy loss of a neural (masked) language model $q_{\boldsymbol{\theta}}(w_i | \boldsymbol{\xi}_i)$, parameterized by $\boldsymbol{\theta}$, that provides distribution over the vocabulary \mathcal{W} given vector representation $\boldsymbol{\xi}_i$ of the W_i ’s context. Then*

$$\mathbb{I}[W_i; \mathbf{x}_i] - \mathbb{I}[W_i; \tilde{\mathbf{x}}_i] \approx \ell(W_i, \tilde{\mathbf{x}}_i) - \ell(W_i, \mathbf{x}_i). \quad (21)$$

30. “you shall know a word by the company it keeps”

31. Just as $\mathbb{I}[T; \boldsymbol{\xi}_i]$ is interpreted as the performance of the best probe that tries to predict the linguistic property T_i given the embedding $\boldsymbol{\xi}_i$ (Pimentel et al., 2020b)

Proof. By Property D.3.1, we have

$$\begin{aligned} \Delta I &:= I[W_i; \mathbf{x}_i] - I[W_i; \tilde{\mathbf{x}}_i] = H[W_i] - H[W_i | \mathbf{x}_i] - H[W_i] + H[W_i | \tilde{\mathbf{x}}_i] \\ &= H[W_i | \tilde{\mathbf{x}}_i] - H[W_i | \mathbf{x}_i] \approx H_q[W_i | \tilde{\mathbf{x}}_i] - H_q[W_i | \mathbf{x}_i], \end{aligned} \quad (22)$$

where $H_q[W_i | \boldsymbol{\xi}_i]$ —a cross entropy—is an estimate of $H[W_i | \boldsymbol{\xi}_i]$ when the true distribution $p(w_i | \boldsymbol{\xi}_i)$ is replaced by a parametric language model $q_{\boldsymbol{\theta}}(w_i | \boldsymbol{\xi}_i)$, which is exactly the cross-entropy loss function of a LM q :

$$H_q[W_i | \boldsymbol{\xi}_i] = \mathbb{E}_{(W_i, C_i) \sim \mathcal{D}}[-\log q_{\boldsymbol{\theta}}(W_i | \boldsymbol{\xi}_i)] =: \ell(W_i, \boldsymbol{\xi}_i). \quad (23)$$

From (22) and (23), we have (21). \square

Remark 3. When we approximate $H[W_i | \mathbf{x}_i] \approx H_q[W_i | \mathbf{x}_i]$, here the parameters $\boldsymbol{\theta}$ of $q = q_{\boldsymbol{\theta}}(w_i | \mathbf{x}_i)$ are only the LM head’s parameters (e.g., the weights and biases of the output softmax layer in BERT), and we assume that \mathbf{x}_i can be *any* representation (e.g., from the last encoding layer of BERT). We assume that $\boldsymbol{\theta}$ are chosen to minimize the KL-divergence between the true distribution $p(w_i | \mathbf{x}_i)$ and the model $q_{\boldsymbol{\theta}}(w_i | \mathbf{x}_i)$, thus making the approximation reasonable.

Similarly, when we approximate $H[W_i | \tilde{\mathbf{x}}_i] \approx H_q[W_i | \tilde{\mathbf{x}}_i]$ we again have the LM head $q_{\boldsymbol{\eta}}(w_i | \tilde{\mathbf{x}}_i)$, whose parameters $\boldsymbol{\eta}$ should be chosen to minimize the KL-divergence between true $p(w_i | \tilde{\mathbf{x}}_i)$ and the model $q_{\boldsymbol{\eta}}(w_i | \tilde{\mathbf{x}}_i)$, for the approximation to be reasonable. In the experimental part (Section 3.3), we do exactly this: after applying INLP (or GRL) to a representation, we fine-tune the softmax layer (but keep the encoding layers frozen) so that it can adapt to the modified representation.

Appendix F. Using GRL in Finetuning Mode

In order to make GRL experiments more comparable to INLP we pretrain BERT (as RoBERTa) on WikiText-103 and finetune it with GRL on the OntoNotes POS annotation. The results for different values of the GRL’s λ are given in Table 6. The majority class

λ	ℓ_{BERT}	Accuracy	λ	ℓ_{BERT}	Accuracy
0	2.062	0.942	10	5.808	0.880
0.1	2.070	0.920	50	6.362	0.844
1	2.647	0.917	100	6.473	0.828
5	3.585	0.913	1000	6.267	0.865

Table 6: Finetuning RoBERTa with GRL on OntoNotes POS annotation.

accuracy on this dataset is 0.126, and we see that no matter how large the λ we take, the GRL still cannot lower the accuracy to this level. This confirms observations of Elazar and Goldberg (2018) that GRL may not always remove all the information from a model it is applied to.

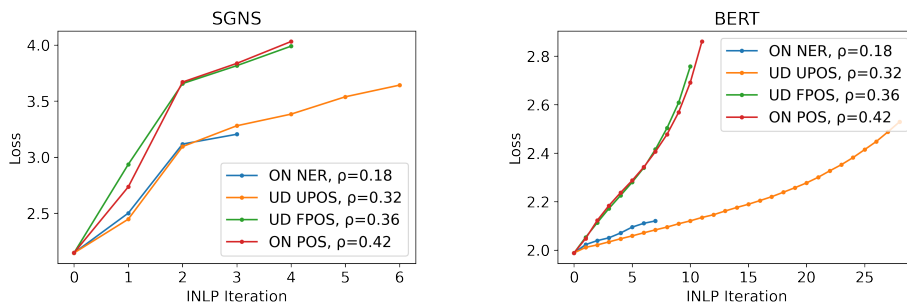


Figure 13: Loss increase w.r.t. INLP iteration for different tasks. ON stands for OntoNotes, UD for Universal Dependencies. The UD EWT dataset has two types of POS annotation: coarse tags (UPOS) and fine-grained tags (FPOS)

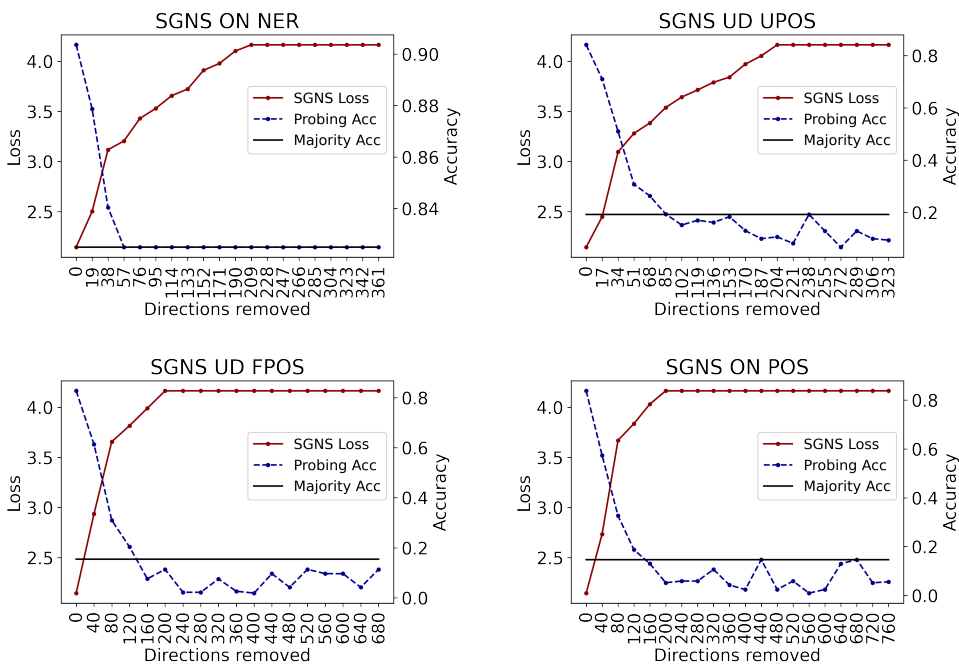


Figure 14: INLP Dynamics for SGNS. ON stands for OntoNotes, UD for Universal Dependencies. The UD EWT dataset has two types of POS annotation: coarse tags (UPOS) and fine-grained tags (FPOS).

Appendix G. INLP Dynamics

Figure 13 shows how quickly and how much the pre-training loss grows when applying the INLP procedure for various linguistic tasks. For each task, we show the minimum number of iterations at which the probing accuracy drops to (or below) the level of the majority accuracy.

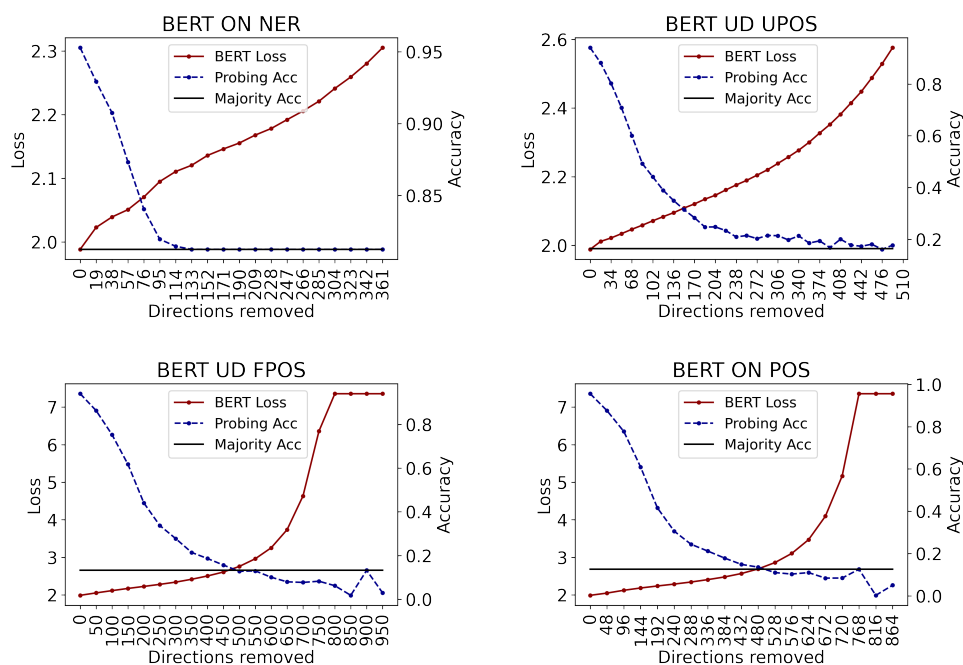


Figure 15: INLP Dynamics for BERT.

We illustrate how pre-training loss and probing accuracy change with respect to INLP iteration in Figures 14 and 15. The number of directions being removed at each INLP iteration is equal to the number of tags in the respective task.

References

- Adi, Y., Kermay, E., Belinkov, Y., Lavi, O., & Goldberg, Y. (2017). Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Allen, C., & Hospedales, T. M. (2019). Analogies explained: Towards understanding word embeddings. In Chaudhuri, K., & Salakhutdinov, R. (Eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Vol. 97 of *Proceedings of Machine Learning Research*, pp. 223–231. PMLR.
- Antoniak, M., & Mimno, D. (2018). Evaluating the stability of embedding-based word similarities. *Trans. Assoc. Comput. Linguistics*, 6, 107–119.
- Arora, S., Li, Y., Liang, Y., Ma, T., & Risteski, A. (2016). A latent variable model approach to pmi-based word embeddings. *Trans. Assoc. Comput. Linguistics*, 4, 385–399.
- Assylbekov, Z. (2020). Sgns implementation in pytorch. <https://github.com/zh3nis/SGNS>. Accessed: 2021-11-07.

- Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In Bengio, Y., & LeCun, Y. (Eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5185–5198, Online. Association for Computational Linguistics.
- Bengio, Y., & Senecal, J. (2003). Quick training of probabilistic neural nets by importance sampling. In Bishop, C. M., & Frey, B. J. (Eds.), *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics, AISTATS 2003, Key West, Florida, USA, January 3-6, 2003*. Society for Artificial Intelligence and Statistics.
- Cettolo, M., Jan, N., Sebastian, S., Bentivogli, L., Cattoni, R., & Federico, M. (2016). The IWSLT 2016 evaluation campaign. In *Proceedings of IWSLT*.
- Chen, T., Frankle, J., Chang, S., Liu, S., Zhang, Y., Wang, Z., & Carbin, M. (2020). The lottery ticket hypothesis for pre-trained BERT networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., & Lin, H. (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Conneau, A., Kruszewski, G., Lample, G., Barrault, L., & Baroni, M. (2018). What you can cram into a single `→` vector: Probing sentence embeddings for linguistic properties. In Gurevych, I., & Miyao, Y. (Eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pp. 2126–2136. Association for Computational Linguistics.
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Mach. Learn.*, 20(3), 273–297.
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., & Solorio, T. (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics.
- Eger, S., Daxenberger, J., & Gurevych, I. (2020). How to probe sentence embeddings in low-resource languages: On structural design choices for probing task evaluation. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pp. 108–118, Online. Association for Computational Linguistics.
- Elazar, Y., & Goldberg, Y. (2018). Adversarial removal of demographic attributes from text data. In Riloff, E., Chiang, D., Hockenmaier, J., & Tsujii, J. (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pp. 11–21. Association for Computational Linguistics.

- Elazar, Y., Ravfogel, S., Jacovi, A., & Goldberg, Y. (2021). Amnesic probing: Behavioral explanation with amnesic counterfactuals. *Trans. Assoc. Comput. Linguistics*, 9, 160–175.
- Ethayarajh, K., Duvenaud, D., & Hirst, G. (2019). Towards understanding linear word analogies. In Korhonen, A., Traum, D. R., & Màrquez, L. (Eds.), *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pp. 3253–3262. Association for Computational Linguistics.
- Ettinger, A., Elgohary, A., & Resnik, P. (2016). Probing for semantic evidence of composition by means of simple classification tasks. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pp. 134–139, Berlin, Germany. Association for Computational Linguistics.
- Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., & Ruppin, E. (2002). Placing search in context: the concept revisited. *ACM Trans. Inf. Syst.*, 20(1), 116–131.
- Firth, J. R. (1957). A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis. Special volume of the Philological Society*, 22(1).
- Frankle, J., & Carbin, M. (2019). The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Ganin, Y., & Lempitsky, V. S. (2015). Unsupervised domain adaptation by backpropagation. In Bach, F. R., & Blei, D. M. (Eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, Vol. 37 of *JMLR Workshop and Conference Proceedings*, pp. 1180–1189. JMLR.org.
- Gittens, A., Achlioptas, D., & Mahoney, M. W. (2017). Skip-gram - zipf + uniform = vector additivity. In Barzilay, R., & Kan, M. (Eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pp. 69–76. Association for Computational Linguistics.
- Gordon, M. A., Duh, K., & Andrews, N. (2020). Compressing BERT: studying the effects of weight pruning on transfer learning. In Gella, S., Welbl, J., Rei, M., Petroni, F., Lewis, P. S. H., Strubell, E., Seo, M. J., & Hajishirzi, H. (Eds.), *Proceedings of the 5th Workshop on Representation Learning for NLP, RepL4NLP@ACL 2020, Online, July 9, 2020*, pp. 143–155. Association for Computational Linguistics.
- Gulordava, K., Bojanowski, P., Grave, E., Linzen, T., & Baroni, M. (2018). Colorless green recurrent networks dream hierarchically. In Walker, M. A., Ji, H., & Stent, A. (Eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pp. 1195–1205. Association for Computational Linguistics.
- Hashimoto, T. B., Alvarez-Melis, D., & Jaakkola, T. S. (2016). Word embeddings as metric recovery in semantic spaces. *Trans. Assoc. Comput. Linguistics*, 4, 273–286.

- Hewitt, J. (2019). structural-probes. <https://github.com/john-hewitt/structural-probes>. Accessed: 2021-11-07.
- Hewitt, J., & Liang, P. (2019). Designing and interpreting probes with control tasks. In Inui, K., Jiang, J., Ng, V., & Wan, X. (Eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 2733–2743. Association for Computational Linguistics.
- Hewitt, J., & Manning, C. D. (2019). A structural probe for finding syntax in word representations. In Burstein, J., Doran, C., & Solorio, T. (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4129–4138. Association for Computational Linguistics.
- Inan, H., Khosravi, K., & Socher, R. (2017). Tying word vectors and word classifiers: A loss framework for language modeling. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Klein, G., Kim, Y., Deng, Y., Senellart, J., & Rush, A. M. (2017). OpenNMT: Open-source toolkit for neural machine translation. In Bansal, M., & Ji, H. (Eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, System Demonstrations*, pp. 67–72. Association for Computational Linguistics.
- Kong, L., de Masson d’Autume, C., Yu, L., Ling, W., Dai, Z., & Yogatama, D. (2020). A mutual information maximization perspective of language representation learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Kuncoro, A., Dyer, C., Hale, J., Yogatama, D., Clark, S., & Blunsom, P. (2018). LSTMs can learn syntax-sensitive dependencies well, but modeling structure makes them better. In Gurevych, I., & Miyao, Y. (Eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pp. 1426–1436. Association for Computational Linguistics.
- Kunz, J., & Kuhlmann, M. (2020). Classifier probes may just learn from linear context features. In Scott, D., Bel, N., & Zong, C. (Eds.), *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pp. 5136–5146. International Committee on Computational Linguistics.
- Lee, J. D., Lei, Q., Saunshi, N., & Zhuo, J. (2020). Predicting what you already know helps: Provable self-supervised learning. *CoRR, abs/2008.01064*.
- Levy, O., & Goldberg, Y. (2014). Neural word embedding as implicit matrix factorization. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., & Weinberger, K. Q. (Eds.), *Advances in Neural Information Processing Systems 27: Annual Conference on*

Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, pp. 2177–2185.

- Linzen, T., Dupoux, E., & Goldberg, Y. (2016). Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Trans. Assoc. Comput. Linguistics*, 4, 521–535.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. *CoRR*, *abs/1907.11692*.
- Mahoney, M. (2011). About the test data. <http://mattmahoney.net/dc/textdata.html>. Accessed: 2021-11-07.
- Marcus, M. P., Santorini, B., & Marcinkiewicz, M. A. (1993). Building a large annotated corpus of english: The penn treebank. *Comput. Linguistics*, 19(2), 313–330.
- Maudslay, R. H., Valvoda, J., Pimentel, T., Williams, A., & Cotterell, R. (2020). A tale of a probe and a parser. In Jurafsky, D., Chai, J., Schluter, N., & Tetreault, J. R. (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 7389–7395. Association for Computational Linguistics.
- McCann, B., Bradbury, J., Xiong, C., & Socher, R. (2017). Learned in translation: Contextualized word vectors. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., & Garnett, R. (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6294–6305.
- Merity, S., Xiong, C., Bradbury, J., & Socher, R. (2017). Pointer sentinel mixture models. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Michael, J., Botha, J. A., & Tenney, I. (2020). Asking without telling: Exploring latent ontologies in contextual representations. In Webber, B., Cohn, T., He, Y., & Liu, Y. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 6792–6812. Association for Computational Linguistics.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013a). Efficient estimation of word representations in vector space. In Bengio, Y., & LeCun, Y. (Eds.), *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. In Burges, C. J. C., Bottou, L., Ghahramani, Z., & Weinberger, K. Q. (Eds.), *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pp. 3111–3119.
- Mikolov, T., Yih, W., & Zweig, G. (2013c). Linguistic regularities in continuous space word representations. In Vanderwende, L., III, H. D., & Kirchhoff, K. (Eds.), *Human Language Technologies: Conference of the North American Chapter of the Association*

of *Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*, pp. 746–751. The Association for Computational Linguistics.

- Mikolov, T., Yih, W., & Zweig, G. (2013d). Linguistic regularities in continuous space word representations. In Vanderwende, L., III, H. D., & Kirchhoff, K. (Eds.), *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*, pp. 746–751. The Association for Computational Linguistics.
- Ney, H., Essen, U., & Kneser, R. (1994). On structuring probabilistic dependences in stochastic language modelling. *Comput. Speech Lang.*, 8(1), 1–38.
- Olah, C. (2015). Visual information theory. <http://colah.github.io/posts/2015-09-Visual-Information/>.
- Ott, M., Baevski, A., Martin, L., & Morton, J. (2019a). Pretraining RoBERTa using your own data. <https://github.com/pytorch/fairseq/blob/main/examples/roberta/README.pretraining.md>. Accessed: 2021-11-07.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., & Auli, M. (2019b). fairseq: A fast, extensible toolkit for sequence modeling. In Ammar, W., Louis, A., & Mostafazadeh, N. (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pp. 48–53. Association for Computational Linguistics.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In Moschitti, A., Pang, B., & Daelemans, W. (Eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pp. 1532–1543. ACL.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In Walker, M. A., Ji, H., & Stent, A. (Eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pp. 2227–2237. Association for Computational Linguistics.
- Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P., Bakhtin, A., Wu, Y., & Miller, A. (2019a). Language models as knowledge bases?. 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)*, pp. 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P. S. H., Bakhtin, A., Wu, Y., & Miller, A. H. (2019b). Language models as knowledge bases?. In Inui, K., Jiang, J., Ng, V., & Wan, X. (Eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural*

- Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 2463–2473. Association for Computational Linguistics.
- Pimentel, T., Saphra, N., Williams, A., & Cotterell, R. (2020a). Pareto probing: Trading off accuracy for complexity. In Webber, B., Cohn, T., He, Y., & Liu, Y. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 3138–3153. Association for Computational Linguistics.
- Pimentel, T., Valvoda, J., Maudslay, R. H., Zmigrod, R., Williams, A., & Cotterell, R. (2020b). Information-theoretic probing for linguistic structure. In Jurafsky, D., Chai, J., Schluter, N., & Tetreault, J. R. (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 4609–4622. Association for Computational Linguistics.
- Prasanna, S., Rogers, A., & Rumshisky, A. (2020). When BERT plays the lottery, all tickets are winning. In Webber, B., Cohn, T., He, Y., & Liu, Y. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 3208–3229. Association for Computational Linguistics.
- Press, W., Teukolsky, S., Vetterling, W., & Flannery, B. (2007). Conditional entropy and mutual information. In *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press.
- Pruksachatkun, Y., Yeres, P., Liu, H., Phang, J., Htut, P. M., Wang, A., Tenney, I., & Bowman, S. R. (2020). jiant: A software toolkit for research on general-purpose text understanding models. In Çelikyilmaz, A., & Wen, T. (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2020, Online, July 5-10, 2020*, pp. 109–117. Association for Computational Linguistics.
- Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: A python natural language processing toolkit for many human languages. In Çelikyilmaz, A., & Wen, T. (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2020, Online, July 5-10, 2020*, pp. 101–108. Association for Computational Linguistics.
- Ravfogel, S., Elazar, Y., Gonen, H., Twiton, M., & Goldberg, Y. (2020). Null it out: Guarding protected attributes by iterative nullspace projection. In Jurafsky, D., Chai, J., Schluter, N., & Tetreault, J. R. (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 7237–7256. Association for Computational Linguistics.
- Ravichander, A., Belinkov, Y., & Hovy, E. H. (2020). Probing the probing paradigm: Does probing accuracy entail task relevance?. *CoRR*, *abs/2005.00719*.
- Reif, E., Yuan, A., Wattenberg, M., Viégas, F. B., Coenen, A., Pearce, A., & Kim, B. (2019). Visualizing and measuring the geometry of BERT. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., & Garnett, R. (Eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Pro-*

- cessing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 8592–8600.
- Rogers, A., Kovaleva, O., & Rumshisky, A. (2020). A primer in bertology: What we know about how BERT works. *CoRR*, *abs/2002.12327*.
- Sahlgren, M., & Carlsson, F. (2021). The singleton fallacy: Why current critiques of language models miss the point. *CoRR*, *abs/2102.04310*.
- Sanh, V., Wolf, T., & Rush, A. M. (2020). Movement pruning: Adaptive sparsity by fine-tuning. *CoRR*, *abs/2005.07683*.
- Saunshi, N., Malladi, S., & Arora, S. (2021). A mathematical exploration of why language models help solve downstream tasks. In *International Conference on Learning Representations*.
- Schlueter, N. (2018). The word analogy testing caveat. In Walker, M. A., Ji, H., & Stent, A. (Eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pp. 242–246. Association for Computational Linguistics.
- Shi, X., Padhi, I., & Knight, K. (2016). Does string-based neural MT learn source syntax?. In Su, J., Carreras, X., & Duh, K. (Eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pp. 1526–1534. The Association for Computational Linguistics.
- Silveira, N., Dozat, T., de Marneffe, M., Bowman, S. R., Connor, M., Bauer, J., & Manning, C. D. (2014). A gold standard dependency corpus for english. In Calzolari, N., Choukri, K., Declerck, T., Loftsson, H., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., & Piperidis, S. (Eds.), *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, pp. 2897–2904. European Language Resources Association (ELRA).
- Stolcke, A. (2002). SRILM - an extensible language modeling toolkit. In Hansen, J. H. L., & Pellom, B. L. (Eds.), *7th International Conference on Spoken Language Processing, ICSLP2002 - INTERSPEECH 2002, Denver, Colorado, USA, September 16-20, 2002*. ISCA.
- Tamkin, A., Singh, T., Giovanardi, D., & Goodman, N. D. (2020). Investigating transferability in pretrained language models. In Cohn, T., He, Y., & Liu, Y. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pp. 1393–1401. Association for Computational Linguistics.
- Tenney, I., Das, D., & Pavlick, E. (2019a). BERT rediscovers the classical NLP pipeline. In Korhonen, A., Traum, D. R., & Màrquez, L. (Eds.), *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28-August 2, 2019, Volume 1: Long Papers*, pp. 4593–4601. Association for Computational Linguistics.
- Tenney, I., Xia, P., Chen, B., Wang, A., Poliak, A., McCoy, R. T., Kim, N., Durme, B. V., Bowman, S. R., Das, D., & Pavlick, E. (2019b). What do you learn from context?

- probing for sentence structure in contextualized word representations. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Tian, R., Okazaki, N., & Inui, K. (2017). The mechanism of additive composition. *Mach. Learn.*, 106(7), 1083–1130.
- Tissier, J., Gravier, C., & Habrard, A. (2019). Near-lossless binarization of word embeddings. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pp. 7104–7111. AAAI Press.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., & Garnett, R. (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008.
- Voita, E., & Titov, I. (2020). Information-theoretic probing with minimum description length. In Webber, B., Cohn, T., He, Y., & Liu, Y. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 183–196. Association for Computational Linguistics.
- Weischedel, R., Pradhan, S., Ramshaw, L., Palmer, M., Xue, N., Marcus, M., Taylor, A., Greenberg, C., Hovy, E., Belvin, R., et al. (2013). Ontonotes release 5.0. LDC2013T19, Philadelphia, Penn.: Linguistic Data Consortium.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., & Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. In Liu, Q., & Schlangen, D. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pp. 38–45. Association for Computational Linguistics.
- Wu, Z., Chen, Y., Kao, B., & Liu, Q. (2020). Perturbed masking: Parameter-free probing for analyzing and interpreting BERT. In Jurafsky, D., Chai, J., Schluter, N., & Tetreault, J. R. (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 4166–4176. Association for Computational Linguistics.
- Yang, Z., Dai, Z., Salakhutdinov, R., & Cohen, W. W. (2018). Breaking the softmax bottleneck: A high-rank RNN language model. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Zhang, Y., Warstadt, A., Li, H., & Bowman, S. R. (2020). When do you need billions of words of pretraining data?. *CoRR*, abs/2011.04946.

- Zhao, M., Lin, T., Mi, F., Jaggi, M., & Schütze, H. (2020). Masking as an efficient alternative to finetuning for pretrained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2226–2241, Online. Association for Computational Linguistics.
- Zhu, Z., & Rudzicz, F. (2020). An information theoretic view on selecting linguistic probes. In Webber, B., Cohn, T., He, Y., & Liu, Y. (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 9251–9262. Association for Computational Linguistics.