
Emergence of Separable Manifolds in Deep Language Representations

Jonathan Mamou^{*1} Hang Le^{*2} Miguel A Del Rio² Cory Stephenson¹
Hanlin Tang¹ Yoon Kim³ SueYeon Chung^{2,4}

Abstract

Deep neural networks (DNNs) have shown much empirical success in solving perceptual tasks across various cognitive modalities. While they are only loosely inspired by the biological brain, recent studies report considerable similarities between representations extracted from task-optimized DNNs and neural populations in the brain. DNNs have subsequently become a popular model class to infer computational principles underlying complex cognitive functions, and in turn, they have also emerged as a natural testbed for applying methods originally developed to probe information in neural populations. In this work, we utilize mean-field theoretic manifold analysis, a recent technique from computational neuroscience that connects geometry of feature representations with linear separability of classes, to analyze language representations from large-scale contextual embedding models. We explore representations from different model families (BERT, RoBERTa, GPT, etc.) and find evidence for emergence of linguistic manifolds across layer depth (e.g., manifolds for part-of-speech tags), especially in ambiguous data (i.e. words with multiple part-of-speech tags, or part-of-speech classes including many words). In addition, we find that the emergence of linear separability in these manifolds is driven by a combined reduction of manifolds' radius, dimensionality and inter-manifold correlations.¹

1. Introduction and Related Work

Many recent studies show notable similarities between representations extracted from task-optimized deep neural networks (DNNs) and neural populations in the brain in sensory systems (Yamins et al., 2014; Khaligh-Razavi & Kriegeskorte, 2014). Computational neuroscience community is increasingly relying on utilizing DNNs as a framework for studying neural correlates underlying complex cognitive functions (Cichy & Kaiser, 2019; Kriegeskorte, 2015). Addressing the question of "how a population of neural units transform representations across multilayered processing stages to implement a cognitive task" is a key challenge in both neuroscience and deep learning. Consequently, developing the techniques to provide insight into neural representation and computation have been an active area of research in both fields (Barrett et al., 2019).

Much prior work on characterizing how information is encoded in DNNs and the brain has focused on the geometric structure underlying the data. In neuroscience, representational similarity analysis (Kriegeskorte & Kievit, 2013) captures the similarity between the stimuli in the geometry of the neural data and deep network representations. Other geometric measures such as geodesics (Hénaff & Simoncelli, 2015), curvature (Hénaff et al., 2019; Fawzi et al., 2018), intrinsic dimension (Ansuini et al., 2019), and canonical correlation analysis (Raghu et al., 2017) have been used to empirically study the complexity of neural population and learned representations in DNNs.

In natural language processing (NLP), recent advances in contextualized word representations such as ELMo (Peters et al., 2018) and BERT (Devlin et al., 2018) have led to significant empirical improvements across many tasks. Concomitant with these advances is an emergent line of work, colorfully referred to as *BERTology*, exploring what aspects of language are being captured by these contextual representations (Zhang & Bowman, 2018; Blevins et al., 2018; Tenney et al., 2019a). One popular approach for analysis is also through the lens of the geometry: Hewitt & Manning (2019) report evidence of a geometric representation of parse trees in embedding from BERT, and Coenen et al. (2019) study the geometric representation of word senses via visualization techniques such as UMAP. Another popu-

^{*}Equal contribution ¹Intel Labs ²Massachusetts Institute of Technology ³Harvard University ⁴Columbia University. Correspondence to: SueYeon Chung <sueyeon@mit.edu>.

Proceedings of the 37th International Conference on Machine Learning, Vienna, Austria, PMLR 119, 2020. Copyright 2020 by the author(s).

¹Code is available at <https://github.com/schung039/contextual-repr-manifolds>.

lar approach for analyzing these representations is through *supervised probes*, i.e. classifiers trained on top of fixed representations to predict certain linguistic properties (e.g., part-of-speech tags, syntactic heads) (Liu et al., 2019a; Tenney et al., 2019b). Supervised probes are conceptually simple and have greatly expanded our understanding of the kinds of linguistic knowledge encoded by these models. However, they are unable to capture the intrinsic geometry underlying the learned representation space, and it is not clear that high accuracy with respect to a probing task necessarily implies that the relevant linguistic structure is being encoded.

In this paper, we apply a recent manifold analysis technique based on replica theory (Chung et al., 2018) that links the geometry of object manifolds to the shattering capacity of a linear classifier as a measure of the amount of information stored about object categories per unit. This method has been used in sensory domains such as visual CNNs (Cohen et al., 2019), visual neuroscience (Chung et al., 2020) and deep speech recognition models (Stephenson et al., 2019) to characterize how object manifolds ‘untangle’ across layers. Here we apply this manifold analysis to study deep language representations, particularly Transformer-based models (Vaswani et al., 2017), for the first time, and show that NLP systems also “untangle” linguistic “objects” relevant for the task.

We present several key findings:

1. Word and linguistic category manifolds emerge across the deep layers of Transformer architectures, in the task-dependent, predictive regime (where the feature vectors are defined on masked tokens), similar to vision and speech deep networks.
2. In word contextualization regime (defined on unmasked tokens), word manifolds strongly decrease in the manifold capacity, becoming less separable across the hierarchy. Linguistic manifolds are affected by the underlying word manifolds, but are counteracted by the contextualization, resulting in linguistic manifolds with a better effective separation compared to word manifolds.
3. The emergence of part-of-speech (POS) manifolds is observed most strongly when the underlying words are ambiguous with multiple POS tags in BERT. POS manifolds further seem to interpolate between word-like geometry and separable contextual geometry, depending on the number of words in each POS class.

In addition, we show the generality of linguistic unangling with word representation manifolds in widely-utilized NLP models. We also show that geometry of fine-tuning learning dynamics can be probed with the tasks congruent, incongruent to the training, to measure the similarity between the

tasks. These results provide geometric evidence for emergence of language representation manifolds, from words to parts-of-speech to named entities, in deep neural networks for natural language processing.

2. Mean-Field Theoretic Manifold Analysis




Figure 1. **Manifold Capacity** measures the linear separability of class manifolds, and it is a generalization of perceptron capacity to classes.

This paper uses mean-field theoretic manifold analysis (Chung et al., 2018; Stephenson et al., 2019) (hereafter, MFTMA technique) as a core analytic tool to measure manifold capacity and other manifold geometric properties (radius, dimension, center correlation) on a subsample of the test dataset.

Given P object manifolds (i.e. feature vectors with their categories) in N feature dimensions, *Manifold capacity* (Fig. 1), $\alpha_M = P/N$, refers to the critical number of object manifolds, P , that can be linearly separated given N features. α_M marks the value above which most manifold dichotomies are inseparable, and below which most are separable with a linear classifier. This is similar to the perceptron capacity for discrete points (Gardner, 1988), except that the unit for the numerator in the perceptron capacity is the number of category manifolds, rather than the number of discrete patterns. The manifold capacity thus measures the linearly separable information about object identity per feature. The manifold capacity can be computed empirically (*Simulation manifold capacity*, α_{SIM} , hereafter), by doing a bisection search to find the critical number of features N such that the fraction of linearly separable random manifold dichotomies is close to $1/2$. MFTMA framework theoretically predicts manifold capacity, and connects it with the geometric properties of category manifolds. As such, MFTMA framework returns four quantities below:

1. **Mean Field Theoretic (MFT) Manifold Capacity** (α_M) estimates the manifold capacity defined above, using the replica mean field formalism introduced by Chung et al. (2018).

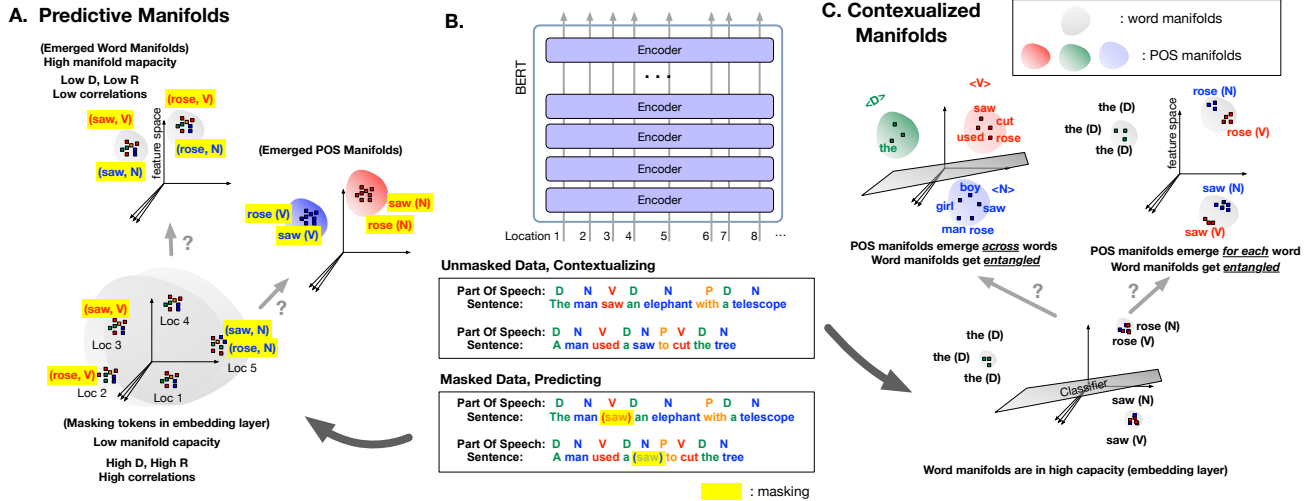


Figure 2. Illustration of a hypothesis that word and linguistic manifolds emerge across a Transformer hierarchy: (A) Predictive manifolds defined on masked tokens. (B) BERT hierarchy. (C) Contextualized manifolds defined on unmasked tokens.

- Manifold Dimension** (D_M) captures the dimensions of an object manifold, and estimates the average embedding dimension of the examples contributing to the decision hyperplane.
- Manifold Radius** (R_M) captures the size of the manifold relevant for linear separability, relative to the norm of the manifold centroid. Small manifold radius implies that the set of examples that determine the decision boundary are tightly grouped.
- Center Correlations** (ρ_{center}) measures the average of absolute values of the pairwise correlations between manifold centroids.

Note that simulation manifold capacity has been reported to be accurately predicted by MFT manifold capacity. We provide their consistency in our data in Fig. 3. In this paper, we use MFT Manifold Capacity (hereafter, Manifold Capacity, α_M). The lower bound of the manifold capacity from data is given by Cover (1965), and reflects the case where there is no manifold structure. Manifold capacity in most random initialized DNNs closely follow lower bound capacity, and we find similar trends in language models (see SM 1.4.1).

The key feature of MFTMA is that manifold capacity can be predicted by the geometric properties of the object manifolds, i.e., Manifold Dimension, Manifold Radius, and their center correlations. Small values for manifold radius, dimension and center correlation result in larger manifold capacity, rendering a more favorable geometry for classification.

3. Experimental Setup

We apply MFTMA to study the geometry of representations from a variety of contextualized word embedding models.

We target Transformer-based models (Vaswani et al., 2017), which compose contextual representations at each layer with self-attention, as they have been shown to produce state-of-the-art results across a large number of NLP tasks. Transformer networks also provide an opportunity to analyze the evolution of representations across layer depth since they typically employ more hidden layers than other neural NLP models.

3.1. Models

BERT-based architectures BERT (Devlin et al., 2018) is a bidirectional Transformer pre-trained using a combination of masked language modeling and next sentence prediction on a large corpus. We also analyzed different architectures derived from BERT: **RoBERTa** (Liu et al., 2019b) modifies key hyperparameters in BERT including removing the next-sentence pre-training objective, and training with much larger mini-batches and learning rates; **ALBERT** (Lan et al., 2019) uses parameter-reduction techniques to lower memory consumption and increase the training speed of BERT; **DistilBERT** (Sanh et al., 2019) is a Transformer model trained by distilling BERT into a smaller model.

GPT architecture **OpenAI GPT** (Radford et al., 2018) is a unidirectional left-to-right Transformer pre-trained using language modeling on a large corpus.

Pretrained models We use the following model versions pretrained on English text: bert-base-cased, albert-base-v1, roberta-base, distilbert-base-uncased, and openai-gpt. All the pre-trained models use a 12-layer transformer except distilbert-base-uncased that uses a 6-layer transformer and have hidden size of 768.

Fine-tuned models In order to analyze how representations change when the model parameters are optimized towards a different task, we also test our approach on models fine-tuned on a downstream task (POS) for different model training *updates steps*.

3.2. Datasets and Manifold Definitions

As noted in section 2, MFTMA begins by assigning each representation to a particular linguistic category (i.e. manifold). We experiment with various word-level categories (derived from common NLP tasks/datasets) that target different language phenomena.

Word A word manifold (*word*) contains several instances of the same word occurring in different contexts. We use the Penn Treebank (PTB) (Marcus et al., 1993) and select 80 word manifolds based on most frequent words in the corpus.

Part-of-speech tags POS (*pos*) tags consists of tags such as *proper nouns (NNP)*, *determiners (DT)*, etc., and are typically considered to target lower-level syntactic phenomena. We select the 33 most frequent tags from PTB.

Semantic tags We use the semantic tagging (*sem-tag*) dataset by Abzianidze & Bos (2017), which annotates words with semantically informative tags. We take the 61 most frequent semantic tags. Some examples of these tags include *comparative positive*, *concept*, *implication*, etc (we refer the reader to the original paper for the full tag definitions). This dataset has also been utilized to analyze contextual word representations in the context of supervised linear probes (Liu et al., 2019a).

Named-entity recognition This task (*ner*) consists of locating and classifying named entity mentioned in text into pre-defined categories such as *person names*, *organizations*, *locations*, etc. It allows for finer-grained manifolds than other tasks since it involves additional segmentation using BIO (*begin*, *inside*, *outside*) tags. We use the tags from the Ontonotes dataset (Weischedel et al., 2011).

Dependency depth All of the above datasets/tags (except words) have been studied by Liu et al. (2019a). Following (Hewitt & Liang, 2019), we study representations stratified by depth in a dependency tree (*dep-depth*). For each contextualized word representation, we use its depth in a dependency tree to assign it to a depth manifold. We select the 22 most frequent depths from PTB.

For each of the word-level tags defined above, we randomly sample 50 word instances per tag to perform the manifold analysis. We average the manifold metrics across five repetitions. For the rest of the paper, we use *linguistic manifolds* to refer to manifolds based on *word*, *pos*, *sem-tag*, *ner* and *dep-depth*. We further use *linguistic category manifolds* to refer to all linguistic manifolds except for depen-

dency depth, since depth in a dependency tree is numeric.

3.3. Feature Extraction

BERT-based models are trained with a masked language modeling objective which randomly replaces words in a sentence with a special [MASK] token. We experiment with two encoding schemes for obtaining the contextualized representations in BERT-based models: *masked* and *unmasked*. In the *masked* case, we use the [MASK] token to obtain the contextualized representation, and assign the representation to the linguistic category of the predicted word (*predictive manifold* in Fig. 2). In the *unmasked* case, we obtain the contextualized word embedding by using the actual token as the input, correspondingly assign the representation to the linguistic category of the input (*contextualized manifold* in Fig. 2). In practice we observe that these two encoding schemes lead to differences in whether linguistic manifolds emerge in earlier or later layers of the network.² If a word is tokenized into multiple tokens (subwords), its word representation is computed as the average of all its subwords' representations.

3.4. Methods

Our mean field theoretic manifold analysis closely follows prior work by Stephenson et al. (2019). We supplement the manifold analysis with two additional techniques.

Distribution of SVM Fields We analyze the distribution of fields (i.e. margins), defined as the signed perpendicular distance between the feature vectors and the optimal linearly separating hyperplane using support vector machines (SVM). We train a slack-SVM with a linear kernel. Given P classes (e.g. P POS tags), we train SVM classifiers for one-versus-rest classification task. The fields are measured only for the feature vectors with positive ground truth label (i.e. "one" of one vs. rest classification). For a given class, the fields from the true positives and false negatives are collected, and normalized by the field distance between the positive and negative class centroids. We collect these normalized fields from all P classes to finally obtain the distribution over fields. We also provide the similar analysis without the normalization (true perpendicular distances to the hyperplane) for reference in SM 1.4.2. The tail of the distribution on the positive side reflects a linear classification accuracy.

PCA visualizations To qualitatively analyze the evolution of representations across layers, we visualize the representations with PCA, where each data point is color-coded according to its tag. When comparing representations across multiple layers, we perform PCA on data across all layers.

²Voita et al. (2019) also analyze representations stratified by masked/unmasked representations, and report considerable differences the evolution of mutual information across layers.

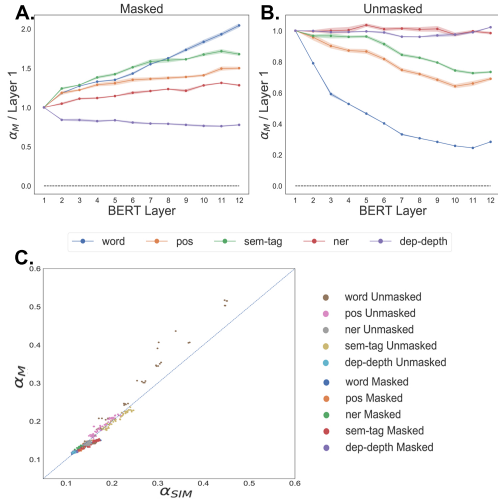


Figure 3. Manifold capacity emerges across various linguistic tags in BERT Manifolds defined by tags for words, POS, sem-tag, NER, dep-depth from (A) masked and (B) unmasked data (Note that empirical capacity matches MFT theory (C)). In A-B, MFT manifold capacity is normalized by manifold capacity of the first layer for each tags ("Capacity/First layer"), to compare the effect across different tasks. Error bars hereafter represent standard error on the mean of the calculated manifold metrics (including 5 random seeds)

4. Results

4.1. Emergence of Separable Language Manifolds

MFTMA analysis We first investigated the manifold capacity of language manifolds in the BERT model, which was trained to predict the word identity of the masked tokens in the output layer, using the datasets described in 3.2. The datasets have two forms: "Masked", to observe the emergent properties of the task-dependent, predictive language manifolds, and "Unmasked", to probe the information content in the contextualized word manifolds.³ The manifold capacity is presented as the relative change compared to capacity of the first layer (Fig. 3, A and B) to enable comparison of the linguistic content embedded in the representations between different datasets.

Since the model is explicitly trained with masked modeling task on the word level, we observe that the word and other linguistic category manifold classes become more separable and increase in their capacity, over the course of the layers, on the masked data (Fig. 3A). Interestingly, other language classes based on POS, ner, sem-tag categories ("linguistic category manifolds") also emerges across the layers, and surprisingly, their relative increase across layers is comparable to the increase in word manifolds, perhaps reflecting the fact that the emergent linguistic category information is used to predict the masked words in the final layer.

³Note that BERT is trained on masked tokens, but uses unmasked tokens at test time.

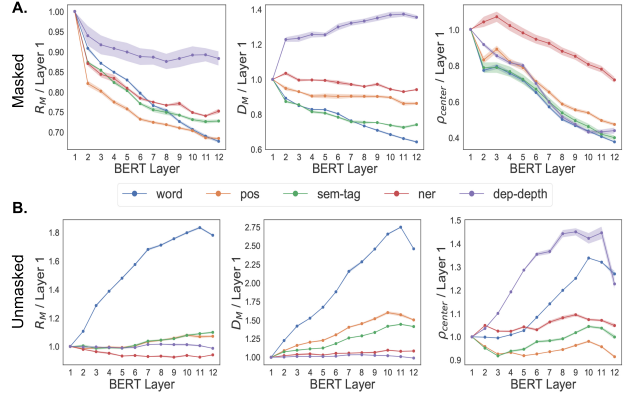


Figure 4. Manifold Dimension, Radius and Center Correlations for various tasks in BERT Geometrical property for different linguistic manifolds: words, POS, sem-tag, ner, dep-depth from masked (A) and unmasked (B) data. These geometrical properties are normalized by the property of the first layer.

On the unmasked data, since the word token is present in the input signal and is dominant, word capacity is much higher than the capacity in the masked data (Fig. 3B). As the input word embedding is already well separated, the word separability is at its highest in the embedding layer as expected. In subsequent layers, these highly separable word manifolds become contextualized and their capacity significantly decreases, as shown in Fig. 3B. For other language manifolds such as POS or ner, where the underlying representations are also based on word features (except with language tags to define partitions based on linguistic categories), the capacity is similarly diminished, although not to the degree compared to the words manifold, due to the effect of the contextualization (Fig. 4). Unlike the masked case, the linguistic manifolds have generally smaller capacity compared to words (Fig. 3B).

While the manifold capacity indicates an emergence of separable manifolds, the mean-field geometric metrics such as the manifold dimension, radius, and correlations tells us "how" the separability arises, i.e., theoretically-grounded geometric evidence of untangling. Across multiple categories (POS, NER, sem-tag, etc.), we find that for the masked data, the increased word and linguistic manifolds capacities are due to the reduction in the manifold radius, dimension, and center correlations (Fig. 4A), similar to prior work in vision and speech (Stephenson et al., 2019). For the unmasked case, the opposite trends are observed (Fig. 4B).

We also find that different linguistic categories show different relative trends. For manifold capacities (Fig. 3), in masked data, words, POS, sem-tag, ner show comparable amount of emergent capacity, while dep-depth shows a general reduction across layers. This might be due to the fact that dep-depth is based on depth of the tree, unlike other category-based tags, and classification-based metrics

Emergence of Separable Manifolds in Deep Language Representations

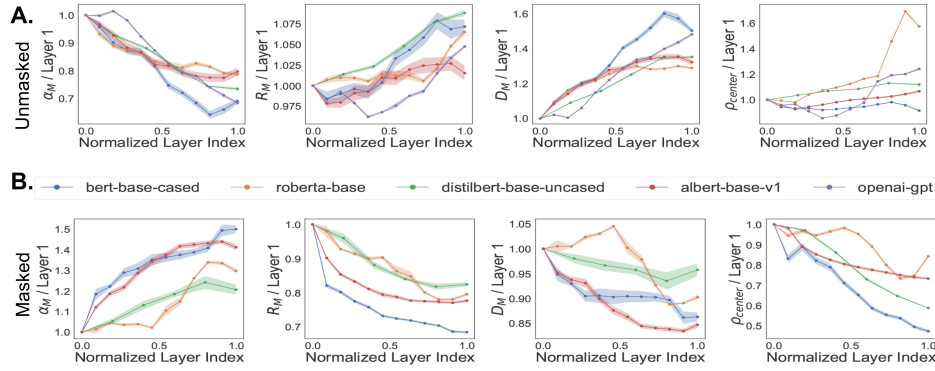


Figure 5. Emergence of language manifolds across different transformer classes, shown by POS manifolds The evolution of manifold capacity, manifold dimension, manifold radius and center correlations within different models (BERT, roBERTa, alBERT, distilBERT and openAI-GPT) for unmasked data (A) and masked data (B). All values are normalized by the first layer. The layer index is normalized so that different models can be plotted together.

might not capture their functional and structural emergence. In unmasked data, a large decrease in the manifold capacity of POS (as compared to ner) potentially indicates that much POS tag information can be derived from words alone, which is consistent with the high accuracy of the most-frequent class baseline for POS tagging.

The trends in the capacity and geometry were similar across different Transformer architectures. The manifold capacities increase across the downstream layers, mediated by the reduction in the manifold dimension, radius, and center correlations in masked data, (Fig. 5), as demonstrated by the POS manifolds. In unmasked data, the capacity is higher overall compared to the capacity in masked data, and the manifold dimension, radius, and correlations are overall smaller, due to the existence of the strong signal in the input. The capacity of unmasked manifolds decrease, and this rate of decrease is steep in Word manifolds, but gradual in linguistic manifolds due to the contextualization effect (see SM 1.4.3 for the rest of the linguistic manifolds across models).

SVM fields statistics analysis We supplement the manifold capacity analysis by analyzing the statistics of the fields, i.e., the signed perpendicular distances between feature vectors and the SVM hyperplane, as described in Sec. 3.4. Fig. 6 shows this for POS manifolds for different train-test splits, and across the BERT model layers. In the masked data (Fig. 6, right column), across all train/test splits, the peak of the fields distribution moves positively away from the origin from early to downstream layers, while the width gets smaller, meaning an increase in the signal to noise ratio, which is also reflected in the measured accuracy (Fig. 6, masked, Inset Figures). In the unmasked data, the peak of the fields distribution moves from the positive side to the origin (in the negative direction), across different train/test split regimes, showing the decreased separability, consistent with the trends observed in the manifold capacity. Interestingly, we find that the trend in linear separability, as measured by

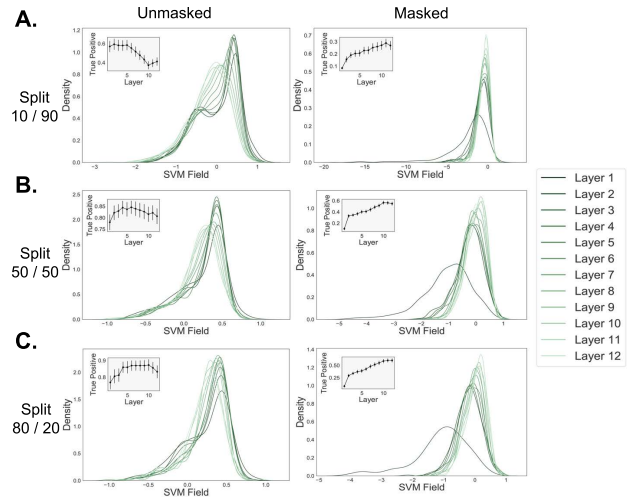


Figure 6. SVM Field Distribution reflects Emergence of linearly separable language manifolds Fields distribution of POS Manifolds on (top left) unmasked and (top right) masked data, on center-normalized data with (Top) 80% Tr/ 20% Test Split, and (Bottom) 10% Tr/ 90% Test Split, on conditioned datasets (true positive + false negative). (Insets in each figure shows true positive rate). 80/20 case is consistent with (Liu et al., 2019a), and 10/90 on unmasked data shows a different trend.

the accuracy across layer depth, is dependent on the size of the training set. With a train/test split of 10/90, the fraction of positive fields (i.e. accuracy) decreases across layers (Fig. 6, Top Left Inset). On the other hand, when we use the same train/test split of 80/20 used by Liu et al. (2019a), we recover their observation that the fraction of positive fields increases across the layers.

4.2. Geometric Evidence for the Stronger Effect of Context in Ambiguous Words

Following our analysis of word and linguistic manifolds with metrics from both manifold theory and SVM fields, we searched for evidence of co-evolution of separable language

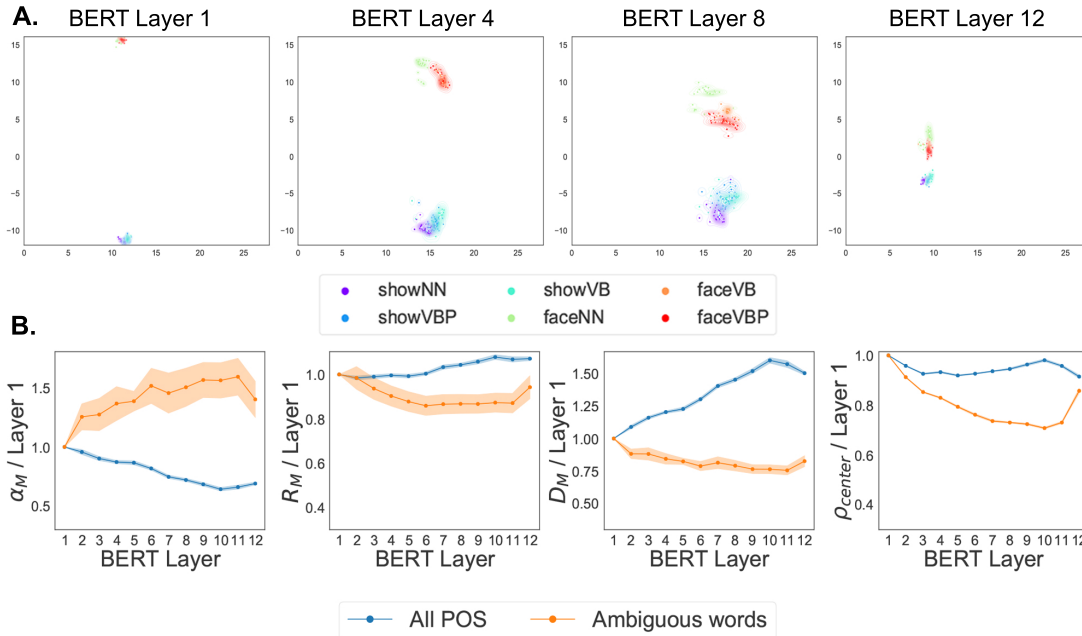


Figure 7. **Visualization and geometric property of POS manifolds with ambiguous words dataset**, unmasked data. (A) PCA visualization of ambiguous words that have multiple possible POS tags across BERT layers. (B) Normalized manifold capacity, radius, dimension, and center correlations of POS manifolds defined on ambiguous words versus on random sampling. For ambiguous words, shaded error bar is calculated from the mean of manifold metrics (including 5 random seeds from the ambiguous word dataset).

manifolds of multiple types by visualizing their representations using PCA. We focus our analysis on two specific data: (1) words that occur across multiple POS tags, which we call "ambiguous words", and (2) open vs. closed POS tags (Lyons, 1977), corresponding to POS tags that have fixed class membership (closed tags) and usually includes only a few words versus those that accept new members (open tags) and usually contains many words.

Geometry of ambiguous words Much of the POS information can be inferred by the choice of a word, regardless of the context, as some words always correspond to a specific POS (e.g., "the", "a", etc.). To test the additional information about POS tags gained by the neighboring context, we focus our analysis here on a specific set of words which occur across at least three different POS tags. This analysis is particularly useful because even for word manifolds, there is an effect of contextualization, observed by the steep decrease in the capacity across layers. Can we reduce the effect of underlying word manifolds by specifically choosing ambiguous words with multiple POS tags? The hypothesis here is that the amount of untangling of POS information might depend on how ambiguous underlying words are.

To test this, we first visualize word embeddings with multiple POS tags with PCA. We find that in the first layer, while the words are highly separated, different POS vectors are embedded very close to each other, and as the repre-

sentations travel downstream, the POS sub-clusters within a same word (different colors) separate, while the overall manifold sizes get larger, clearly showing the competing effect of contextualization. This is a known challenge in the "untangling" deep networks, observed in vision; that is, the network needs to embed the whole data in high dimensional space, while separating and compressing the categorical data in low dimensional space, all using the same network parameters (Cohen et al., 2019; Recanatesi et al., 2019).

With the competing effect of contextualization on manifold geometry observed by PCA visualization, we then quantify these effects with our manifold analysis metric. One question is: while it's clear that the POS information is segregating conditioned on each word (Fig. 7, Top), are POS manifolds defined across many words also segregating, and do they segregate more, if the underlying words are ambiguous? Our results suggest that indeed BERT layers untangle POS manifolds from ambiguous words more than they untangle typical POS manifolds. In Fig. 7B, while POS manifold capacity of randomly sampled dataset decreases (blue line), POS manifold capacity of manifold sampled from ambiguous words increase across layers (orange line). Furthermore, we also observe that in the ambiguous dataset, manifold radius, dimension and center correlation also decrease across layers, suggesting a trend of untangling (Cohen et al., 2019; Stephenson et al., 2019), as opposed to the entangling trend in randomly sampled POS dataset.

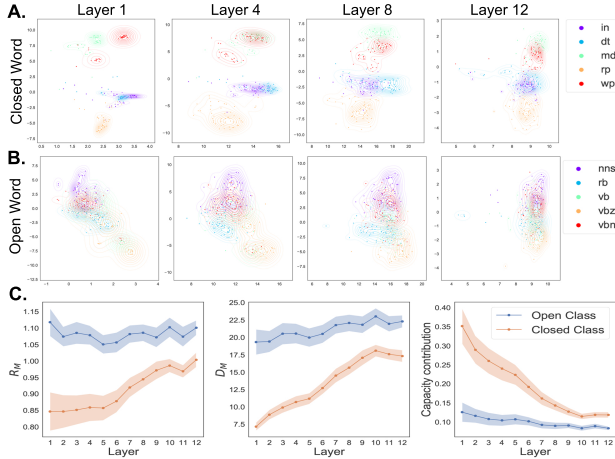


Figure 8. Linguistic manifold visualization and geometric property of closed and open word class PCA visualization for 5 randomly selected POS classes from open word class (A) and closed word class (B) (density is estimated by bivariate kernel density estimate). (C) Manifold radius, dimension and capacity contribution of POS classes in open word class and closed word class across BERT layer. Error bar represents standard error between different POS classes within open word or closed word class.

Interestingly, the untangling effect is strongest in the middle layers, which is consistent with the PCA visualization (Fig.7). This effect maybe the result of the model objective to predict the word identity in the last layer, making several POS manifolds within a word becomes more entangled.

We observe that this analysis can explain the ostensible contradiction between the decrease in manifold capacity across layer depth that we measured and the increase in accuracy across layer depth that has generally been observed from supervised probes (Liu et al., 2019a). In particular, the visualizations indicate that there is an *overall* entangling of representations in later layers (when averaged across all words), leading to decrease in manifold capacity. However, there is an untangling of representations *within* each word, contributing to higher probe accuracy in later layers.

Geometry of open vs. closed POS tags In addition to ambiguous words, we analyze POS classes with different number of words within each POS class, in unmasked data. We compare the geometry of closed class POS categories, corresponding to a small number of distinct words, and open class POS categories, with a large number of distinct words. Do closed-class POS manifolds show a similar geometrical transformation properties as word manifolds? Open-class manifolds, as embedded across many words, might be highly entangled in the input layer; do they show more emergent separability compared to closed-class manifolds?

Based on their PCA visualizations (Fig. 8A-B)), closed word POS classes are indeed already well separated in the

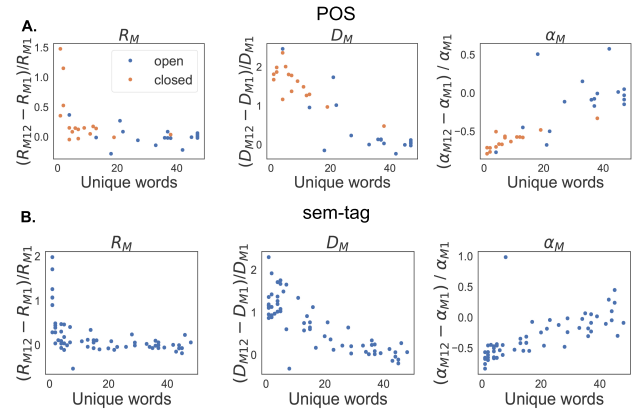


Figure 9. Relationship between manifold geometrical evolution and number of unique words within each manifold relationship between manifold radius, dimension and capacity ratio to the number of unique word in each manifold. The ratio is calculated by the formula $(X_{12} - X_1)/X_1$ with X_1 and X_{12} be the measurement from layer 1 (first layer) and layer 12 (last layer) respectively.

embedding layer, and over the layers, these POS manifolds become more tangled, similar to the word manifolds trends seen previously (Fig. 3B). On the other hand, open POS classes are quite entangled in the beginning embedding layer, and their change in separability appears to be relatively small (Fig. 8B). Applying our manifold capacity method clearly shows that the measured manifold dimension and radii increases across layers in the closed class, but changes minimally for the open class (Fig. 8C). Similarly, the capacity contribution⁴ decreases significantly for closed class, but changes minimally for the open class.

Finally, we study the relationship between manifold geometry of output/input ratio (defined in Fig. 9 caption) and number of unique words within a linguistic manifold for POS and sem-tag tasks (Fig. 9). While manifold capacity ratio is correlated with the number of unique words, manifold radius and dimension ratio is anti-correlated. This result implies that for POS and sem-tag, the linguistic manifolds with large number of words counteract the effect of expanding underlying word manifolds the most, signifying the structural evidence of contextualization on untangling of POS and sem-tag representations.

4.3. Geometry of Learning Dynamics and Task-transferability

In addition to evaluating the MFTMA metrics on the pre-trained networks, the analysis can also be done as the training progresses. Here, we trained a popular fine-tuning

⁴Manifold capacity is defined as $\alpha = \langle \alpha_\mu^{-1} \rangle_\mu^{-1}$, where α_μ is a capacity contribution from the μ th manifold. For capacity contribution for open (closed) classes (Fig. 8C), μ refers to manifold indices corresponding to open (closed) classes.

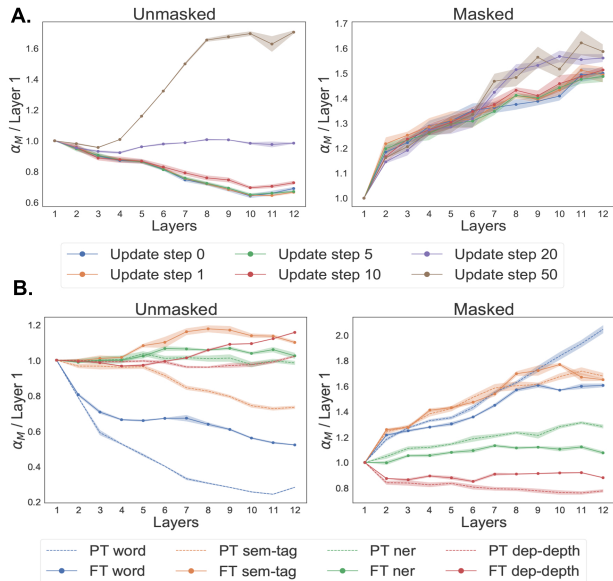


Figure 10. Geometric evidence of fine tuning and task transfer representations (A) Emergence of POS manifolds on (left) unmasked, (right) masked data across update steps of POS fine-tuning on BERT. (B) Linguistic manifolds on (left) unmasked, (right) masked data on (solid) final step of POS fine-tuning on BERT, compared with pretrained BERT (dashed).

task, POS, on a pre-trained BERT, and trace the geometry of learned representations over the course of the training. Fig. 10A shows the evolution of representation geometry in different stages of learning epochs, measured by the capacity of manifolds measured on the same tags as the training task (POS). Across training steps, POS manifolds in contextualizing regime (unmasked data) become more separable across layers, resulting in the regime where the manifold capacity increases across layers in the final step. Surprisingly, POS fine-tuning has little effect on masked data.

Furthermore, we characterize how these representations fine-tuned with POS task transfer to a different task, by measuring the manifold capacity with other linguistic tags, i.e., word, ner, sem-tag, and dep-depth. In unmasked data, the linguistic manifolds increase in their overall capacities with POS fine-tuning, showing the evidence of the task transfer. In masked data, most linguistic manifolds also improve in their capacity by a small amount, with an exception of word and NER manifolds, where an overall entanglement is observed between pretrained and fine-tuned models. Note that masked token is never seen during fine-tuning.

In addition, we supply the analysis on the manifold capacity vs. task performance (F1, precision, recall) across the fine-tuning updates in SM 1.4.4.

5. Discussion

In this paper, we studied the emergent geometric properties of word and linguistic object manifolds and their linear separability, as measured by the shattering manifold capacity. Across different network models and datasets, we find that language manifolds emerge across the layers of the hierarchy. Particularly in the predictive manifolds defined on the masked data, the manifold capacity consistently increases, similar to the ‘untangling’ phenomena observed in visual and auditory sensory systems (Cohen et al., 2019; Chung et al., 2020; Stephenson et al., 2019). Contextualized manifolds defined on unmasked words also show implicit emergence of separable linguistic geometry, counteracting the strong entangling effect of word manifolds. Interestingly, the untangling effect of contextualization is stronger in words and linguistic tags that are ambiguous.

The emergence of increasing manifold capacity (accompanied by the emergent manifold geometry) on the masked data is reflective of the fact that the representations emerge across layers to become more “similar” to the last layer’s output, consistent with the prior reports on the increasing mutual information between the intermediate and the last layers of contextual embedding models (Voita et al., 2019). The decrease of the word manifold capacity on unmasked data implies that information about the input word generally gets lost as representations get transformed downstream, similar to the reduced mutual information between input and intermediate layers (Voita et al., 2019). Furthermore, our metric goes beyond the constraints of the comparative measure, as the manifold capacity measures the amount of object information for any categories, allowing analysis of higher level linguistic categories such as POS.

Contextualization is a competition between gaining information from neighboring words, without losing information about the original word. Just as much as the original word representation is enhanced by the context, the same word representation is used to enhance the representation of other words. Measuring how distributed systems such as Transformers balance this multitude of information flows and implement linguistic information in a mixed representation is a theoretical challenge. Enabled by the recent manifold analysis technique, we report the quantifiable structural evidence of the evolution of language manifolds in connection to their linear separability in widely used language models.

Our methodology and results suggest many interesting future directions. We hope that our work will motivate: (1) the theory-driven geometric analysis of emergent representations underlying complex tasks such as prediction and contextualization; (2) the development of new theoretical frameworks that link the representation geometry and tasks with underlying causal structures; (3) the future study of language representations in the brain via the lens of geometry.

Acknowledgements

We thank Haim Sompolsky, Larry Abbott, Ev Fedorenko, Roger Levy, Greg Wayne, Jon Gauthier and Abhinav Ganesh for helpful discussions. The work was funded by Intel Research Grant, NSF NeuroNex Award DBI-1707398, and The Gatsby Charitable Foundation.

References

- Abzianidze, L. and Bos, J. Towards universal semantic tagging. *arXiv preprint arXiv:1709.10381*, 2017.
- Ansuini, A., Laio, A., Macke, J. H., and Zoccolan, D. Intrinsic dimension of data representations in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 6109–6119, 2019.
- Barrett, D. G., Morcos, A. S., and Macke, J. H. Analyzing biological and artificial neural networks: challenges with opportunities for synergy? *Current opinion in neurobiology*, 55:55–64, 2019.
- Blevins, T., Levy, O., and Zettlemoyer, L. Deep rnns encode soft hierarchical syntax. *arXiv preprint arXiv:1805.04218*, 2018.
- Chung, S., Lee, D. D., and Sompolsky, H. Classification and geometry of general perceptual manifolds. *Physical Review X*, 8(3):031003, 2018.
- Chung, S., Dapello, J., Cohen, U., DiCarlo, J., and Sompolsky, H. Separable manifold geometry in macaque ventral stream and dcnns. In *Computational and Systems Neuroscience (Cosyne)*, 2020.
- Cichy, R. M. and Kaiser, D. Deep neural networks as scientific models. *Trends in cognitive sciences*, 2019.
- Coenen, A., Reif, E., Yuan, A., Kim, B., Pearce, A., Viegas, F., and Wattenberg, M. Visualizing and Measuring the Geometry of BERT. In *Proceedings of NeurIPS*, 2019.
- Cohen, U., Chung, S., Lee, D. D., and Sompolsky, H. Separability and geometry of object manifolds in deep neural networks. *bioRxiv*, pp. 644658, 2019.
- Cover, T. M. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE transactions on electronic computers*, (3):326–334, 1965.
- Devlin, J., Chang, M.-W., Lee, K., , and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. abs/1810.04805, 2018. URL <http://arxiv.org/abs/1810.04805>.
- Fawzi, A., Moosavi-Dezfooli, S.-M., Frossard, P., and Soatto, S. Empirical study of the topology and geometry of deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3762–3770, 2018.
- Gardner, E. J. The space of interactions in neural network models. *Journal of Physics A: Mathematical and General*, 21(1):257–270, 1988. ISSN 0305-4470. doi: 10.1088/0305-4470/21/1/030. URL <http://stacks.iop.org/0305-4470/21/i=1/a=030>.
- Hénaff, O. J. and Simoncelli, E. P. Geodesics of learned representations. *arXiv preprint arXiv:1511.06394*, 2015.
- Hénaff, O. J., Goris, R. L., and Simoncelli, E. P. Perceptual straightening of natural videos. *Nature neuroscience*, pp. 1, 2019.
- Hewitt, J. and Liang, P. Designing and Interpreting Probes with Control Tasks. In *Proceedings of EMNLP*, 2019.
- Hewitt, J. and Manning, C. D. A Structural Probe for Finding Syntax in Word Representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4129–4138, 2019.
- Khaligh-Razavi, S.-M. and Kriegeskorte, N. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11), 2014.
- Kriegeskorte, N. Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual review of vision science*, 1:417–446, 2015.
- Kriegeskorte, N. and Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends in cognitive sciences*, 17(8):401–412, 2013.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. Albert: A lite bert for self-supervised learning of language representations, 2019.
- Liu, N. F., Gardner, M., Belinkov, Y., Peters, M. E., and Smith, N. A. Linguistic knowledge and transferability of contextual representations. *CoRR*, abs/1903.08855, 2019a. URL <http://arxiv.org/abs/1903.08855>.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized bert pretraining approach, 2019b.

- Lyons, J. *Semantics* (vol. 1 & vol. 2), 1977.
- Marcus, M. P., Santorini, B., and Marcinkiewicz, M. A. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2): 313–330, 1993. URL <https://www.aclweb.org/anthology/J93-2004>.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 2227–2237, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1202. URL <https://www.aclweb.org/anthology/N18-1202>.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf, 2018.
- Raghu, M., Gilmer, J., Yosinski, J., and Sohl-Dickstein, J. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Advances in Neural Information Processing Systems*, pp. 6076–6085, 2017.
- Recanatani, S., Farrell, M., Advani, M., Moore, T., Lajoie, G., and Shea-Brown, E. Dimensionality compression and expansion in deep neural networks. 2019.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 2019.
- Stephenson, C., Feather, J., Padhy, S., Elibol, O., Tang, H., McDermott, J., and Chung, S. Untangling in invariant speech recognition. In *Advances in Neural Information Processing Systems*, pp. 14368–14378, 2019.
- Tenney, I., Das, D., and Pavlick, E. Bert rediscovers the classical nlp pipeline. *arXiv preprint arXiv:1905.05950*, 2019a.
- Tenney, I., Xia, P., Chen, B., Wang, A., Poliak, A., McCoy, R. T., Kim, N., Van Durme, B., Bowman, S. R., Das, D., et al. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*, 2019b.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention Is All You Need. In *Proceedings of NeurIPS*, 2017.
- Voita, E., Sennrich, R., and Titov, I. The Bottom-up Evolution of Representations in the Transformer: A Study with Machine Translation and Language Modeling Objectives. In *Proceedings of EMNLP*, 2019.
- Weischedel, R., Hovy, E., Marcus, M., Palmer, M., Belvin, R., Pradhan, S., Ramshaw, L., and Xue, N. Ontonotes: A large training corpus for enhanced processing. *Handbook of Natural Language Processing and Machine Translation*. Springer, pp. 59, 2011.
- Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014.
- Zhang, K. W. and Bowman, S. R. Language modeling teaches you more syntax than translation does: Lessons learned through auxiliary task analysis. *arXiv preprint arXiv:1809.10040*, 2018.