

The Stem Cell Hypothesis: Dilemma behind Multi-Task Learning with Transformer Encoders

Han He

Department of Computer Science
Emory University
Atlanta, GA 30322 USA
han.he@emory.edu

Jinho D. Choi

Department of Computer Science
Emory University
Atlanta, GA 30322 USA
jinho.choi@emory.edu

Abstract

Multi-task learning with transformer encoders (MTL) has emerged as a powerful technique to improve performance on closely-related tasks for both accuracy and efficiency while a question still remains whether or not it would perform as well on tasks that are distinct in nature. We first present MTL results on five NLP tasks, POS, NER, DEP, CON, and SRL, and depict its deficiency over single-task learning. We then conduct an extensive pruning analysis to show that a certain set of attention heads get claimed by most tasks during MTL, who interfere with one another to fine-tune those heads for their own objectives. Based on this finding, we propose the *Stem Cell Hypothesis* to reveal the existence of attention heads naturally talented for many tasks that cannot be jointly trained to create adequate embeddings for all of those tasks. Finally, we design novel parameter-free probes to justify our hypothesis and demonstrate how attention heads are transformed across the five tasks during MTL through label analysis.

1 Introduction

Transformer encoders (TEs) have established recent state-of-the-art results on many core NLP tasks (He and Choi, 2019; Yu et al., 2020; Zhang et al., 2020). However, their architectures can be viewed “over-parameterized” as downstream tasks may not need all those parameters, prone to cause an overhead in computation. One promising approach to mitigate this overhead is multi-task learning (MTL) where a TE is shared across multiple tasks; thus, it needs to be run only once to generate final embeddings for all tasks (Clark et al., 2019b).

Despite the success in MTL on closely-related tasks such as language understanding (Wang et al., 2018) or relation extraction (Chen et al., 2020; Lin et al., 2020), MTL on core NLP tasks (e.g., tagging, parsing, labeling) whose decoders are very distinct has not been well-studied. This work employs the state-of-the-art decoders on five core tasks for MTL

and thoroughly analyzes interactions among those tasks to explore a possibility of reducing the computation overhead from TEs. Surprisingly, our experiments depict that models jointly trained by MTL give lower accuracy than ones trained individually, that is against findings from previous work. In fact, models jointly trained with all five tasks perform the worst among any other combination (Section 3).

These experimental results urge us to figure out why MTL on core tasks with a shared TE leads to worse performance than its single-task counterparts. Our exploration begins by detecting essential heads for each task by forcing the TE to use as few attention heads as possible while maintaining accuracy similar to a fully-utilized encoder. Our experiments reveal that all five tasks rely on almost the same set of attention heads. Hence, they compete for those heads during MTL, causing to blur out features extracted by individual tasks. Thus, we propose the *Stem Cell Hypothesis*, likening these talented attention heads to *stem cells*, which cannot be fine-tuned for multiple tasks that are very distinct (Section 4).

To validate this hypothesis, many parameter-free probes are designed to observe how every attention head is updated while trained individually or jointly. Intriguingly, we find that heads not fine-tuned for any task can still give remarkably high performance to predict certain linguistic structures, confirming the existence of *stem cells* inherently more talented; it is consistent with previous work stating that TEs carry on a good amount of syntactic and semantic knowledge (Tenney et al., 2019; Liu et al., 2019a; Jawahar et al., 2019; Hewitt and Manning, 2019). After single-task learning, probing results typically improve along with the task performance, illustrating that the *stem cells* are developed into more task-specific experts. On the contrary, MTL often drops both probing and task performance, supporting our hypothesis that attention heads lose expertise when exposed to multiple teaching signals that may conflict to one another (Section 5).

The *Stem Cell Hypothesis* is proposed to shed light on a possible direction to MTL research using TEs, comprising an unbearable amount of parameters, by wisely assigning attention heads to downstream tasks. Although most analysis in this study is based on BERT, we also provide extensive experimental results and visualization of other recent TEs including RoBERTa (Liu et al., 2019c), ELECTRA (Clark et al., 2020) and DeBERTa (He et al., 2020) in §A.4 to further demonstrate the generality of our hypothesis. To the best of knowledge, this is the first time that a comprehensive analysis of attention heads is made for MTL on those core tasks by introducing novel parameter-free probing methods.¹

2 Related Work

A small portion of our work overlaps with multi-task learning. MTL with pre-trained transformers specifically in NLP (Wang et al., 2018; Clark et al., 2019b; Liu et al., 2019b; Kondratyuk and Straka, 2019; Chen et al., 2020; Lin et al., 2020) has been widely studied. Most work focus on neural architecture design to encourage beneficial message passing across tasks. Our MTL framework adopts conventional architecture and applies tricks of batch sampling (Wang et al., 2019) and loss balancing.

Most of our work falls into the analysis of BERT, especially from a linguistic view. Since BERT was introduced, studies on explaining why BERT works have never stopped. The most related studies are those trying to study the linguistic structures learnt by BERT. Among them, Tenney et al. (2019) and Liu et al. (2019a) showed part-of-speech, syntactic chunks and roles can be discovered from BERT embeddings. Using a supervised probe, Hewitt and Manning (2019) successfully discover full dependency parse trees. The encoded dependency structure is also supported by Jawahar et al. (2019) using probes on embeddings. Apart from these parameterized probes, parameter-free approaches (Clark et al., 2019a; Wu et al., 2020) also agree with the existence of rich linguistic knowledge in BERT, which is closely related to our probing methods.

What remains unclear is the impact of fine-tuning on TEs. Using supervised probes, Peters et al. (2019) claim that fine-tuning adapts BERT embeddings to downstream tasks, which is later challenged by Hewitt and Liang (2019) since su-

pervised probe itself can encode knowledge. Then, Zhao and Bethard (2020) propose a methodology to test such encoding of a linguistic phenomenon by comparing the probing performance before and after fine-tuning. Our probing methods align with these unsupervised probes while focus more on explaining the impact of multi-task learning.

3 Multi-Task Learning

Our goal of MTL is to build a joint model sharing the same encoder but using a distinct decoder for each task that outperforms its single-task counterparts while being faster and more memory efficient. Our model adapts hard parameter sharing (Caruana, 1993) such that all decoders take the same hidden states generated by the shared encoder as input and make task-specific predictions in parallel.

3.1 Shared Encoder

For main experiments, BERT (Devlin et al., 2019) is used as the shared encoder although our approach can be adapted to any transformer encoders (§A.4). Every token gets split into subtokens by BERT; eventually, the average of the last layer’s hidden states generated for those subtokens is used as the final embedding of that token. Additionally, word dropout is applied for generalization by replacing random subtokens with [MASK] during training.

3.2 Task-Specific Decoders

Five tasks are experimented, part-of-speech tagging (POS), named entity recognition (NER), dependency parsing (DEP), constituency parsing (CON), and semantic role labeling (SRL). For each task, a state-of-the-art decoder is adopted (except for POS) to provide a modern benchmark for MTL on these tasks, and simplified to build an efficient model.

POS A linear layer is used as a POS decoder that takes the final embedding of each token from BERT and generates the output vector where each dimension gives the score of a particular POS tag.

NER The biaffine decoder (Yu et al., 2020) is used for NER. For simplification, document context, fastText and character-level embeddings as well as variational BiLSTM encoding from the original approach are removed.

DEP The biaffine decoder (Dozat and Manning, 2017) is used for DEP as well. For simplification, part-of-speech tags, character-level embeddings and the variational BiLSTM are removed in our

¹All our resources including source codes and models are public available at <https://github.com/emorynlp/stem-cell-hypothesis>.

	POS	NER	DEP	CON	SRL	MTL-5
POS	98.32 ± 0.02	98.28 ± 0.01	98.28 ± 0.02	98.30 ± 0.02	98.27 ± 0.02	98.25 ± 0.01
NER	89.34 ± 0.24	89.04 ± 0.14	89.43 ± 0.14	88.38 ± 0.06	89.18 ± 0.25	88.94 ± 0.10
DEP	94.04 ± 0.02	94.06 ± 0.07	94.24 ± 0.03	94.12 ± 0.04	94.12 ± 0.03	93.84 ± 0.08
CON	94.23 ± 0.03	94.38 ± 0.04	94.33 ± 0.02	94.43 ± 0.03	94.25 ± 0.05	94.10 ± 0.05
SRL	82.92 ± 0.10	82.39 ± 0.07	82.05 ± 0.03	83.17 ± 0.02	82.93 ± 0.08	82.30 ± 0.13

Table 1: Performance of single-task learning (main diagonal highlighted in gray), multi-task learning on all 5 tasks (MTL-5), and multi-task learning on every pair of the tasks (non-diagonal cells; e.g., DEP’s row in NER’s column is the DEP result of the joint model between DEP and NER). See also Table 12 for similar results of other TEs.

approach. Also, the final embedding of [CLS] from BERT is used to represent the root node.

CON The two-stage CRF decoder is used for CON (Zhang et al., 2020). The unlabeled bracket scorer is optimized using a tree-structure CRF objective on unlabeled constituents. The encoding layer from the original approach is substituted by BERT. Also, [CLS] and [SEP] in BERT are used to represent [BOS] and [EOS], respectively.

SRL The end-to-end span ranking decoder (He et al., 2018) is used for SRL. The attention-based span representations are replaced by the averaged embeddings as suggested by Xia et al. (2019). For simplification, a linear layer is used as the ranker instead of the biaffine one since they have shown similar performance in our preliminary experiments.

3.3 Data and Loss Balancing

During multi-task training, batches from different tasks are shuffled together and randomly sampled to optimize the shared encoder and the corresponding decoder. Following Wang et al. (2019), a task is sampled based on a probability proportional to its dataset size raised to the power of 0.8. To balance the losses of all tasks, a running average of every task is monitored and its loss is updated as follow:

$$\mathcal{L}'_t = \frac{\sum_{\forall i} \bar{\mathcal{L}}_i}{\bar{\mathcal{L}}_t} \cdot \mathcal{L}_t$$

\mathcal{L}_t is the current loss of the task t , $\bar{\mathcal{L}}_t$ is the running average of the most recent 5 losses of t , and \mathcal{L}'_t is the updated loss of t . This balancing method normalizes the loss of each task to the same magnitude and has shown to prevent MTL from being biased to specific tasks in our preliminary experiments.

3.4 MTL Experiments

Our models are experimented on the OntoNotes 5 (Weischedel et al., 2013) using the data split suggested by Pradhan et al. (2013). Table 1 illustrates performance of all models using the following evaluation metrics - POS: accuracy, NER: span-level

labeled F1, DEP: labeled attachment score, CON: constituent-level labeled F1, SRL: micro-averaged F1 of (predicate, argument, label). Every model is trained 3 times and their average score and standard deviation on the test set are reported. For DEP, the gold trees from CON are converted into the Stanford dependencies v3.3.0 (de Marneffe and Manning, 2008). Detailed descriptions about the experimental settings are provided in Appendix A.2.

Single-task learning models are first trained then compared to the MTL model trained on all 5 tasks (MTL-5). Interestingly, MTL-5 is outperformed by its single-task counterparts for all tasks. Due to the high complexity of MTL-5, it is hard to tell which combinations of tasks introduce negative transfer. Thus, we conduct MTL on every pair of the tasks to observe if there is any task combination that yields a positive result (non-diagonal cells in Table 1).

Among the 10 pairwise task combinations, none of them derives a win-win situation. NER results are generally improved with MTL although results on the other tasks are degraded, implying that NER takes advantage of the other tasks by hurting their performance. SRL is also benefited from CON although it is not the case for the other way around. Results of other recent TEs reveal similar patterns as shown in Appendix A.4.

4 Pruning Analysis

To answer why MTL leads to suboptimal results in Section 3, pruning strategies are applied to BERT such that only attention heads absolutely necessary to get the best performance are kept for every task. This allows us to see if there exists a common set of heads that multiple tasks want to claim and train for only their objectives, which can cause conflicts for those heads to be shared across all tasks.

4.1 Pruning based on L_0 -Regularization

BERT is essentially a stack of multi-head attention layers and there is a wide consensus that different layers learn distinct knowledge (Lin et al., 2019;

	Performance				% of Attention Heads Kept			PS/S	
	STL	STL-SP	STL-DP	MTL-DP	STL-SP	STL-DP	MTL-DP	STL	STL-DP
POS	98.32±0.02	98.22±0.03	98.35±0.02	98.28±0.01	53.24±4.07	40.51±1.61	50.70±1.20	405	1,245
NER	89.04±0.14	88.87±0.10	89.05±0.08	88.78±0.13	57.87±2.63	49.77±7.13	50.70±1.20	661	700
DEP	94.24±0.03	94.08±0.10	94.22±0.06	93.92±0.06	63.66±5.02	50.00±2.08	50.70±1.20	241	601
CON	94.43±0.03	94.16±0.08	94.24±0.03	94.16±0.05	45.37±0.40	44.91±0.40	50.70±1.20	191	397
SRL	82.93±0.08	83.01±0.05	83.11±0.16	82.77±0.10	82.41±5.99	53.24±4.07	50.70±1.20	299	326

Table 2: Results of single-task learning (STL), STL with static pruning (STL-SP), STL with dynamic pruning (STL-DP), and multi-task learning on the 5 tasks with dynamic pruning (MTL-DP). PS/S: processed samples per second for speed comparison. The STL Performance column is equivalent to the main diagonal in Table 1. See also Table 13 for pruning results of other recent TEs.

Hewitt and Manning, 2019; Jawahar et al., 2019; Liu et al., 2019a; Tenney et al., 2019). Inspired by this, we analyze if each head learns features unique to process different tasks. First, L_0 -regularization is applied (Louizos et al., 2018) to encourage BERT to use as few heads as possible during training. In particular, a binary variable z_j is assigned to the j 'th head and multiplied to the output of that head (see Vaswani et al. (2017) for $\mathbf{Q}, \mathbf{K}, \mathbf{V}, d_k$):

$$\text{Attention}^{(j)}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = z_j \cdot \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V}$$

Unfortunately, these binary variables $\mathbf{z} = \{z_j : \forall_{j=[1,\ell]}\}$ (ℓ : total # of heads) are discrete and non-differentiable so cannot be directly learnt using gradient based optimization. To allow for efficient continuous optimization, each \mathbf{z}_j is then relaxed as a random variable drawn independently from a continuous random distribution. Specifically, the relaxed \mathbf{z} is re-parameterized by its inverse of the cumulative density function (CDF) as $G_\alpha(\mathbf{u})$. It is sampled as follows, where α is a learnable parameter of the inverse CDF, U is the uniform distribution over the interval $[0, 1]$ and $\mathbf{u} = \{u_j : \forall_{j=[1,\ell]}\}$ denotes the iid samples from it:

$$\mathbf{u} \sim U(0, 1) \Rightarrow \mathbf{z} = G_\alpha(\mathbf{u})$$

Then, the *Hard Concrete Distribution* (Louizos et al., 2018) is chosen for \mathbf{z} , which gives the following form of $G_\alpha(\mathbf{u})$ that is differentiable, where (l, r) defines the interval that $g_\alpha(\mathbf{u})$ can be stretched into ($l < 0, r > 1$):

$$g_\alpha(\mathbf{u}) = \text{sigmoid}(\log \mathbf{u} - \log(\mathbf{1} - \mathbf{u}) + \alpha)$$

$$G_\alpha(\mathbf{u}) = \min(\mathbf{1}, \max(\mathbf{0}, g_\alpha(\mathbf{u}) \times (r - l) + l))$$

By sampling \mathbf{u} and applying the Monte Carlo approximation, the learnable L_0 -objective is obtained in a closed form, which gets jointly optimized with

a task specific loss or the balanced MTL loss:

$$\mathbb{E}_{\mathbf{u} \sim U(0,1)} [\mathbf{z}] = \text{sigmoid}\left(\alpha - \log \frac{-l}{r}\right)$$

$$\mathbb{E}_{\mathbf{u} \sim U(0,1)} [L_0] = \sum_{j=1}^n \mathbb{E}[z_j] \quad (1)$$

4.2 Pruning Strategies

Two types of pruning strategies, static and dynamic, are applied for the attention head analysis:


Static Pruning We refer to the conventional two-stage train-then-prune as static pruning (SP) since it fine-tunes the encoder first then freezes the decoder for pruning (Voita et al., 2019).

Dynamic Pruning Since SP requires twice the efforts to obtain a pruned model, we propose a new method that simultaneously fine-tunes and prunes. This strategy is referred to as dynamic pruning (DP) since the decoder dynamically adapts to the encoder that is being pruned during training, as opposed to SP which instead freezes the decoder. DP is found to be more effective in our experiments.

All pruning models are trained for 3 runs with different random seeds and the best checkpoints by scores on development sets are kept. Once trained, $\mathbb{E}_{\mathbf{u} \sim U(0,1)} [\mathbf{z}] \in (0, 1)$ is used as a measure of how much each head is being utilized.

4.3 Pruning Experiments

Table 2 shows single-task learning (STL) results using SP and DP on the 5 tasks. Our DP strategy consistently performs better than the SP strategy as it shows higher accuracy on all tasks and prunes significantly greater numbers of heads except for CON. Compared to the STL models without any pruning, the STL-DP models perform well or slightly better for POS/NER/SRL due to the L_0 -regularization, yet use $\approx 50\%$ fewer numbers of heads.



(f) 3-run utilization of the MTL-DP model, where each run is encoded in a RGB channel. Darker indicates higher utilization.

(g) Average head utilization rates among the 5 tasks in 3 runs. Darker cells indicate higher utilization rates.

Figure 1: Head utilization of the STL-DP models (a - e, g) and the MTL-DP model (f). The detailed quantification of correlations is shown in Table 3. See also Figure 9, 11 and 13 for similar visualization of other recent TEs.

Comparing the STL-DP models across different tasks, SRL requires more heads than DEP and CON, which require more than POS. This aligns with the intuition behind their difficulty levels as *semantic* > *syntactic* > *lexical* relations. On the other hand, NER requires more heads than CON because of the world knowledge it needs to capture from the data; thus, this knowledge is more scattered.

Finally, DP is applied to MTL-5 (MTL-DP), which shows slightly higher accuracy than MTL-5 in Table 1 except for NER by pruning 50% of the heads. This might imply that all tasks want to claim a similar set of heads even though about a half of the heads are underutilized during MTL training.

4.4 Pruning Visualization

To visualize the utilization of heads across runs, the utilization rate $z_{j,t}^{(r)}$ of the j 'th head in the r 'th run for the task t is encoded to the RGB channels:

$$R/G/B_{j,t} = 255 \times (1 - z_{j,t}^{(1/2/3)})$$

For instance, RGB of (0, 0, 0) is *black* indicating that the head is 100% utilized in all 3 runs. Based

on this scheme, the head utilizations of all STL-DP models as well as the MTL-DP model are plotted in Figures 1a ~ 1f. To depict the overlaps of utilized heads across tasks, z 's are averaged over all STL-DP models for all runs then plotted as a grayscale heatmap (Figure 1g) using the following scheme:

$$H_j = \frac{1}{15} \sum_{r=1}^3 \sum_{t=1}^5 z_{j,t}^{(r)}$$

Consistent head utilization by runs As shown in Figures 1a ~ 1e and the main diagonal in Table 3, the head utilization per task seems quite similar across different runs, especially for syntactic/semantic tasks such as DEP/CON/SRL. The head utilization of POS seems to be random because it is a simple task so that high performance can be achieved by a small set of the re-utilized heads. This consistency across different runs is an essential prerequisite for the following analyses.

Consistent head utilization across tasks Figures 1a ~ 1e show that all of these tasks are mostly utilizing heads from layers 5 to 8 (looking like M).

	POS	NER	DEP	CON	SRL	MTL-5
POS	75.51	74.77	76.45	83.54	73.32	78.34
NER	74.77	73.28	73.57	75.23	74.44	66.99
DEP	76.45	73.57	89.89	89.27	91.20	84.39
CON	83.54	75.23	89.27	83.99	83.90	80.71
SRL	73.32	74.44	91.20	83.90	81.01	80.07
MTL-5	78.34	66.99	84.39	80.71	80.07	85.88

Table 3: Adjusted R-squared of 3-run head utilization rates using the third run as the dependent variable (main diagonal highlighted in gray) and Pearson Correlation Coefficient of averaged head utilization rates between each pair of models (non-diagonal cells).

In contrast to [Jawahar et al. \(2019\)](#) and [Tenney et al. \(2019\)](#), our findings suggest that the middle layers also provide rich surface and semantic features, which are aligned with [Liu et al. \(2019a\)](#) showing that both POS and chunking tasks perform the best when heads from the middle layers are utilized.

Consistent head utilization by STL and MTL

Figures 1f and 1g illustrate almost the identical utilization patterns, implying that the MTL-DP model re-uses a very similar set of heads used by the STL-DP models. According to [Vaswani et al. \(2017\)](#), the representation capacity of every head is limited by the design of multi-head attention. Since (1) a similar set of heads are used across multiple tasks and (2) the limited representation capacity of individual heads confines them to only specific tasks, forcing them for MTL leads to worse results. Given this analogy, we propose the following hypothesis:

There exists a subset of attention heads in a transformer called “stem cells” that are commonly used by many tasks, which cannot be jointly trained for multiple tasks that are very different in nature.

We refer to this claim as the *Stem Cell Hypothesis* and seek to test it through the probing analysis.

5 Probing Analysis

This paper hypothesizes the existence of *stem cells*, which cannot be trained to create adequate embeddings to be shared by multiple tasks that are not so similar. This section provides empirical evidence to this hypothesis by probing what roles each attention head plays once fine-tuned for end tasks.

5.1 Probing Methods

Previous studies on probing transformer encoders have focused on layer-level analysis limited to supervised probing ([Hewitt and Manning, 2019](#); [Tenney et al., 2019](#); [Lin et al., 2019](#); [Jawahar et al.,](#)

[2019](#); [Zhao and Bethard, 2020](#)). This section introduces probes on the head-level instead to analyze the impact of fine-tuning on every individual head. Since developing supervised probes on hundreds of heads requires extensive resource, parameter-free probing methods are used in this study.

Attention Probes Attention between two words often matches a certain linguistic relation that gives a good indicator to knowledge encoded in the head. Our decoders for DEP and SRL learn relationships between head/dependent words and predicate/argument words respectively, which can be directly benefited from these attentions. Thus, the attention matrix from each head is used as the probe of that head. Following [Clark et al. \(2019a\)](#), an undirected edge is created between each word and its most attending word. First, the subtoken-subtoken attention matrix is converted into a word-word matrix by averaging the attention probabilities of each multi-subtoken word. The $\arg \max$ of each row r in the attention matrix is then calculated, denoted as \mathbf{g}_r , and evaluated on the basis of each task.

For DEP, directions of the gold arcs are removed and compared against the predicted arcs as follows ($h|d$: the index of a head|dependent word, (h, d) : an undirected arc from the gold tree, n : # of arcs):

$$\frac{1}{n} \sum_{\forall (h,d)} \mathbb{1}(\mathbf{g}_h = d \parallel \mathbf{g}_d = h)$$

For SRL, we design a new probing method to evaluate how each word in the argument span is attended to the head word in its predicate (p : the index of a predicate head word, T^p : word indices in the span of p ’s argument, m : # of predicate-argument pairs):

$$\frac{1}{m} \sum_{\forall p, \forall T^p} \mathbb{1}(\mathbf{g}_p \in T^p \parallel \exists t \in T^p : \mathbf{g}_t = p)$$

Only head words in the predicates are used for this analysis, which affects verb-particle constructions (e.g., only *throw* is used for *throw away*). Moreover, not all words in an argument span are necessarily important to add meaning to its predicate. We will explore these aspects in the future.

Attended-Value Probes POS/NER/CON can be viewed as tasks to find and label spans in a sentence, where the span is a word for POS, a sequence of consecutive words for NER and CON, where a span can be overlapped with another span for CON. For

these tasks, we again present a new probing method, depicted in Algorithm 1, that predicts the label of each span based on its representation:

Algorithm 1: Attended-Value Probing

Function PseudoCluster (\mathbf{H}, \mathcal{S}):

```

 $\mathbf{C} \leftarrow \mathbf{0} \in \mathbb{R}^{m \times d}, \mathbf{n} \leftarrow \mathbf{0} \in \mathbb{R}^{m \times 1}$ 
foreach  $(b, e, \ell) \in \mathcal{S}$  do
     $\mathbf{C}_\ell \leftarrow \mathbf{C}_\ell + \text{mean}(\mathbf{H}_b, \dots, \mathbf{H}_e)$ 
     $\mathbf{n}_\ell \leftarrow \mathbf{n}_\ell + 1$ 
return  $\mathbf{C}/\mathbf{n}$ 

Function AVProbe ( $\mathbf{C}, \mathbf{H}$ ):
    return  $\arg \max(\text{cossim}(\mathbf{C}, \mathbf{H}))$ 

```

The attended-value matrix $\mathbf{H} \in \mathbb{R}^{n \times d}$ is created by multiplying the attention matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ to the value matrix $\mathbf{V} \in \mathbb{R}^{n \times d}$ (Section 4.1) such that $\mathbf{H} = \mathbf{AV}$ (n : sentence length, d : embedding size, abbreviated d_k). \mathcal{S} is the set of gold spans, m is the total number of labels, (b, e, ℓ) denotes the indices of the beginning word, the ending word, and the label respectively, and `cossim` is a cosine similarity function with broadcasting enabled.

With Algorithm 1, the centroid of each label is obtained through `PseudoCluster` then used to predict labels of all spans. Note that for `CON`, only constituents on the height-3 (right above the `POS`-level) are used for this analysis. We experimented with constituents on higher levels, which did not show good correlation with model performance as the spans got longer and noisy. We plan to design another probing method for deeper analysis in `CON`.

5.2 Probing Experiments

Probing experiments are conducted on all attention heads in the pre-trained BERT (Devlin et al., 2019) and fine-tuned models trained by single-task learning (STL; diagonal in Table 1), pairwise multi-task learning (MTL; other cells in Table 1), and 5-task multitask learning (MTL-5; last column in Table 1) using the two probing methods, attention probes and attended-value probes (Section 5.1). For each model, the head with the highest probing accuracy among 144 heads (12 heads per layer, 12 layers) is selected per label. Since every model is developed 3 times using different random seeds for better generalization (Section 3.4), 3 heads are selected per label, which are averaged to get the final probing score for that label. The full probing results with respect to all labels are described in Appendix A.3.

Two important observations are found from these experiments. (1) Even without fine-tuning, certain heads perform remarkably well on particular labels, confirming the existence of *stem cells* (Sec. 5.2.1). (2) Most heads show higher performance once fine-tuned; nonetheless, MTL does not always enhance them for all tasks. In fact, MTL models show improvement on only a few labels (Sec. 5.2.2), while MTL-5 models show no benefit for most labels.²

5.2.1 Pluripotent Stem Cells

The probing results of pre-trained attention heads from BERT (before fine-tuning) are visualized to verify the existence and pluripotency of *stem cells*. These results show very high accuracies for many labels, confirming the existence of *stem cells*. As they reside in the same pre-trained model, their pluripotency is therefore implied. Specifically, the number of probing tasks is 203 (`POS`:49, `NER`:19, `SRL`:67, `DEP`:45, `CON`:23 as shown in Appendix A.3.3) which is larger than the number of heads (144) in BERT-base by itself. Not all of them provide task specific knowledge as shown in our pruning experiments (Section 4), so the number of utilized heads is even smaller. As a result, some heads must play multiple roles in different tasks.

Dependency Parsing For `DEP`, probing results from the best performing heads with respect to their layers for all labels are plotted in Figure 2, some of which are even comparable to supervised results.




Figure 2: DEP layer analysis of pre-trained heads.

The best performing head of BERT finds the `ROOT` of a sentence with a 96.25% accuracy without any supervision, demonstrating its ability to convey the concept. Furthermore, the identification of `ROOT`

²The probing results for `DEP` and `SRL` are omitted here due to the space limit but explained in Appendix A.3.

happens mostly at the early stage of inference, i.e. in layers 2 and 3. This finding may conflict with the idea of syntactic features getting learned in middle layers (Jawahar et al., 2019). It takes the argument from Tenney et al. (2019) a step further suggesting that syntax can be encoded in early layers of TEs.

Semantic Role Labeling As shown in Figure 3, probing shows promising results on many semantic roles. Specifically, numbered arguments (ARG0–4) are recognized in layers 5 to 7, while modifiers are identified in layers 8 to 10 with > 80% accuracies, including ARGM-MOD (modals), ARGM-DIR (directional), ARGM-EXT (extent), ARGM-LVB (light-verb), and ARGM-COM (comitative). Unlike DEP that most labels are learned within the first 7 layers, SRL requires 7+ layers to be learned such that no role reaches the peak before layer 5. This implies that semantic roles take more efforts to be learned than syntactic dependencies.




Figure 3: SRL layer analysis of pre-trained heads.

5.2.2 Stem Cells Specialization

Though *stem cells* are pluripotent, they develop into specialized ones in STL and lose specialities in MTL according to the following comparisons of best performing heads across BERT, STL and MTL models.

Part-of-Speech Tagging Figure 4 compares the heads in the STL model against the other models; the y -axis shows the probing results from the model in the x -axis subtracted by the results of the STL model. Labels (sorted by frequency) with negative scores for BERT imply that STL performs better on those labels than BERT (without getting fine-tuned), whereas negative labels with the other models (e.g., NER, DEP) imply that the joint models

perform worse than the STL model on those labels.

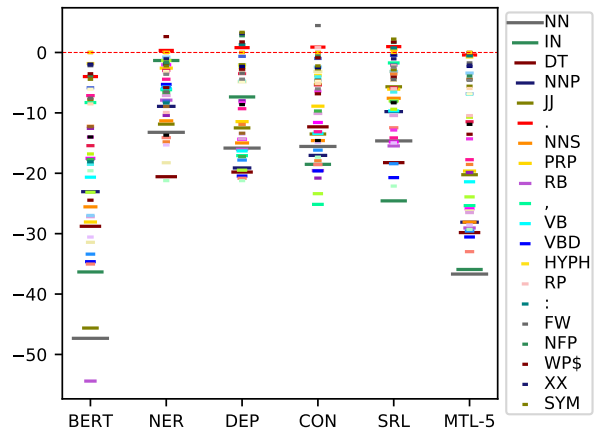


Figure 4: POS probing results comparison.

For POS, MTL degrades performance for most labels compared to STL. Even without getting fine-tuned, the pre-trained BERT model performs very well on punctuation labels, which is expected. The performance on WP\$ (possessive wh-pronoun) is significantly improved with NER, DEP, and SRL as a possessive wh-pronoun (e.g., *whose*) often follows a name or is used in a relative clause that plays an important role in DEP and SRL.

Named Entity Recognition For NER, BERT detects PERCENT, MONEY, LAW, LANGUAGE, NORP (national|religious|political groups) and PRODUCT with over 90% probing accuracy, probably due to the rich set of those entities present in pre-training data. Although most joint models degrade probing results for nearly every entity type, POS and DEP improve upon more entity types than the other tasks (Figure 5), which is consistent with the results illustrated in Table 1.




Figure 5: NER probing results comparison.

Constituency Parsing As shown in Figure 6, POS improves the most number of constituent types but also causes the largest drop among the MTL models for the most frequent type, NP (noun phrase). This contributes to related constituent types such as RB (adverb) in ADVP (RB phrase), WRB (wh-adverb) in WHADVP (WRB phrase), and UH (interjection) in INTJ (UH phrase). Its dramatic decrease on the NP performance might be due to the internal lexical complexity in NP.




Figure 6: CON probing results comparison.

Regarding NER, its boost on ADVP can be due to temporal entities (TIME, DATE) nested within ADVP such as (ADVP (NP *one year ago*)), where *one year ago* is a DATE entity. As for PP (preposition phrase), it usually follows the induction rule of $PP \rightarrow IN/TO + NP$, where the NP is often an named entity (e.g., (PP (TO *to*) (NP *Mary*))). Regarding DEP, it mainly improves wh-phrases like WHNP and WHADVP, which correspond to *nsubj* and *advmod* dependency relations, respectively. Regarding SRL, it slightly improves NML (nominal modifiers) and FRAG (fragment), which may be ascribed to the strength of the span-based SRL not requiring constituency structures for decoding.

Note that in these probing analyses, we selected the best performing heads from each model independently as their locations are not regular anymore. Without the pruning objective (Equation 1), the locations of the best performing heads are non-regular possibly due to the knowledge transfer between *stem cells* and *non-stem cells*. Thus, knowledge transferring from *stem cells* to *non-stem cells* becomes much easier when the models are free to use as many heads as they want. In fact, when fine-tuned without the pruning objective, many *stem cell* attention heads transfer their knowledge to *non-stem cell* heads to get specialized. It is a phe-

nomenon frequently observed in many previous works (Tenney et al., 2019; Liu et al., 2019a; Jawahar et al., 2019) and this work (Section 5.2.2) that certain layers achieve best performance of certain tasks. Given that the *stem cells* of BERT are mostly in middle layers (Section 4.4), we believe that the best performing layers or heads in lower or higher layers are the results of transfer learning on *stem cells*. In reality, a *stem cell* also moves from its original area (e.g., bone marrow) to another area (e.g., bone surface) to get specialized.

6 Conclusion

This study analyzes interference on the 5 core tasks by highlighting naturally talented attention heads, whose importance turns out to be invariant for many downstream tasks. The *Stem Cell Hypothesis* states that these talented heads are like *stem cells* that can develop into experts but not all-rounders. Our hypothesis is validated by several novel parameter-free probes, revealing the interfered representations of *stem cells*. We will adapt this work to more tasks and languages for broader generality in the future.

References

- Richard A. Caruana. 1993. Multitask Learning: A Knowledge-Based Source of Inductive Bias. In *Proceedings of the 10th International Conference on Machine Learning, ICML'93*, pages 41–48.
- Daoyuan Chen, Yaliang Li, Kai Lei, and Ying Shen. 2020. Relabel the Noise: Joint Extraction of Entities and Relations via Cooperative Multiagents. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5940–5950, Online. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019a. What Does BERT Look at? An Analysis of BERT’s Attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Urvashi Khandelwal, Christopher D. Manning, and Quoc V. Le. 2019b. BAM! Born-Again Multi-Task Networks for Natural Language Understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5931–5937, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *ICLR*.

- Marie-Catherine de Marneffe and Christopher D. Manning. 2008. [The Stanford Typed Dependencies Representation](#). In *Coling 2008: Proceedings of the workshop on Cross-Framework and Cross-Domain Parser Evaluation*, pages 1–8, Manchester, UK. Coling 2008 Organizing Committee.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep Biaffine Attention for Neural Dependency Parsing](#). In *Proceedings of the 5th International Conference on Learning Representations, ICLR'17*.
- Han He and Jinho D. Choi. 2019. [Establishing Strong Baselines for the New Decade: Sequence Tagging, Syntactic and Semantic Parsing with BERT](#). In *Proceedings of the 33rd International Florida Artificial Intelligence Research Society Conference, FLAIRS'19*.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018. [Jointly Predicting Predicates and Arguments in Neural Semantic Role Labeling](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 364–369, Melbourne, Australia. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- John Hewitt and Percy Liang. 2019. [Designing and Interpreting Probes with Control Tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A Structural Probe for Finding Syntax in Word Representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What Does BERT Learn about the Structure of Language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Dan Kondratyuk and Milan Straka. 2019. [75 Languages, 1 Model: Parsing Universal Dependencies Universally](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2779–2795, Hong Kong, China. Association for Computational Linguistics.
- Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. [A Joint Neural Model for Information Extraction with Global Features](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7999–8009, Online. Association for Computational Linguistics.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. [Open Sesame: Getting inside BERT’s Linguistic Knowledge](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019a. [Linguistic Knowledge and Transferability of Contextual Representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019b. [Multi-Task Deep Neural Networks for Natural Language Understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Christos Louizos, Max Welling, and Diederik P. Kingma. 2018. [Learning Sparse Neural Networks through \$L_0\$ Regularization](#). In *International Conference on Learning Representations*.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To Tune or Not to Tune? Adapting Pretrained Representations to Diverse Tasks](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (ReL4NLP-2019)*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. [Towards Robust Linguistic Analysis using OntoNotes](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152,

- Sofia, Bulgaria. Association for Computational Linguistics.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT Rediscovered the Classical NLP Pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019. [Can You Tell Me How to Get Past Sesame Street? Sentence-Level Pretraining Beyond Language Modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4465–4476, Florence, Italy. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nanwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium, Philadelphia, PA*.
- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020. [Perturbed Masking: Parameter-free Probing for Analyzing and Interpreting BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4166–4176, Online. Association for Computational Linguistics.
- Qingrong Xia, Zhenghua Li, and Min Zhang. 2019. [A Syntax-aware Multi-task Learning Framework for Chinese Semantic Role Labeling](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing* (*EMNLP-IJCNLP*), pages 5382–5392, Hong Kong, China. Association for Computational Linguistics.
- Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. [Named Entity Recognition as Dependency Parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6470–6476, Online. Association for Computational Linguistics.
- Yu Zhang, Houquan Zhou, and Zhenghua Li. 2020. [Fast and Accurate Neural CRF Constituency Parsing](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4046–4053. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Yiyun Zhao and Steven Bethard. 2020. [How does BERT’s attention change when you fine-tune? An analysis methodology and a case study in negation scope](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4729–4747, Online. Association for Computational Linguistics.

A Appendix

A.1 Corpus Statistics

Table 4 and 5 describes statistics of the POS, NER, DEP, CON, SRL datasets used in our experiments.

	Sentences	Tokens
TRN	75,187	1,299,312
DEV	9,603	163,104
TST	9,479	169,579

Table 4: POS, DEP, CON, SRL statistics of OntoNotes 5. TRN/DEV/TST: training/development/evaluation set.

	Sentences	Tokens
TRN	59,924	1,088,503
DEV	8,528	147,724
TST	8,262	152,728

Table 5: NER statistics of OntoNotes 5. TRN/DEV/TST: training/development/evaluation set.

A.2 Hyper-Parameter Configuration

The hyper-parameters used in our models are described in Table 6.

BERT Encoder	
name	bert-base-cased
layers tag	12
hidden size	768
subword dropout	0.2
Adam Optimizer	
encoder lr	5e-5
decoder lr	1e-3
ϵ	1e-8
epochs	30
warm up	10
NER Decoder	
MLP size	150
dropout	0.5
DEP Decoder	
arc MLP size	500
rel MLP size	100
dropout	0.33
CON Decoder	
span MLP size	500
label MLP size	100
dropout	0.33
SRL Decoder	
argument ratio	0.8
predicate ratio	0.4
span width size	20
max Span width	30
label MLP size	100
dropout	0.2

Table 6: Hyper-parameters settings.

A.3 Extra BERT Probing Results

A.3.1 Dependency Parsing




Figure 7: DEP probing results comparison.

Compared to STL, a similar trending with POS and NER can be observed that MTL improves only certain relations as shown in Figure 7. Among the 4 tasks, POS improves the least tags in terms of both overall accuracy and probing accuracy. It improves `dobj` (direct object) and `expl` (expletive) possibly due to that its decoder needs to assign a verb tag to the `ROOT` verb and `EX` (existential there) to “there”, enhancing the representations of these two. Regarding NER, it mainly improves modifiers that modify nouns which comprises named entities. In the case of CON, modifiers and complement arguments are improved, most of which usually reside in NP (Noun Phrase) or VP (Verb Phrase) phrases, placing upper bounds on the distance of dependencies. As regards SRL, it improves subjects and clausal relations.

A.3.2 Semantic Role Labeling




Figure 8: SRL probing results comparison.

In comparison to STL illustrated in Figure 8, POS and CON mainly improves ARG0-3 (agent, patient, instrument, benefactive, attribute and starting point) and some modifiers including ARGM-TMP (temporal), ARGM-CAU (cause), ARGM-PRD (secondary predication), ARGM-EXT (extent), ARGM-PNC (purpose) and ARGM-REC (reciprocals). Both POS and CON reveal syntactic functions which appear to coordinate attentions on semantic roles in a similar way. Regarding NER and DEP, they improve arguments that include referent or pronouns (R-ARG1, R-ARG0) and modifiers (ARGM-NEG which is negation, ARGM-CAU, ARGM-PRD, ARGM-EXT), possibly due to the biaffine decoders they employ analogously enhance the heads.

A.3.3 Detailed Probing Results

With STL as the baseline, probing results for each task are recorded in Table 7 to 11 respectively. In each table, the first column shows the labels sorted by their frequencies in the test set, the second column shows the mean probing results over 3 runs for STL which is the baseline, other columns show the probing results for BERT and MTL models with their differences against the STL baseline (in parentheses, green indicates improvements and red indicates decreases). BERT probing scores are usually much lower than STL, which is expected as it is not fine-tuned on the task data. MTL probing scores are also generally lower than STL, implying interference with other joint tasks.

	STL	BERT	NER	DEP	CON	SRL	MTL-5
NN	91.44	44.11(-47.33)	78.22(-13.22)	75.61(-15.83)	75.88(-15.56)	76.80(-14.64)	54.75(-36.69)
IN	89.79	53.45(-36.34)	88.46(-1.33)	82.43(-7.36)	71.27(-18.52)	65.22(-24.57)	53.86(-35.93)
DT	93.93	65.15(-28.78)	73.36(-20.57)	74.13(-19.80)	81.62(-12.31)	75.69(-18.24)	64.11(-29.82)
NNP	90.21	67.16(-23.05)	81.29(-8.92)	71.09(-19.12)	73.17(-17.04)	80.42(-9.79)	62.11(-28.10)
JJ	83.59	37.96(-45.63)	71.73(-11.86)	71.10(-12.49)	70.09(-13.50)	77.91(-5.68)	63.35(-20.24)
.	95.02	91.03(-3.99)	95.36(+0.34)	95.82(+0.80)	95.91(+0.89)	95.99(+0.97)	94.61(-0.41)
NNS	91.75	66.19(-25.56)	80.43(-11.32)	76.77(-14.98)	77.16(-14.59)	84.17(-7.58)	63.66(-28.09)
PRP	93.76	65.69(-28.07)	91.17(-2.59)	82.31(-11.45)	84.87(-8.89)	87.68(-6.08)	74.13(-19.63)
RB	83.33	28.92(-54.41)	75.44(-7.89)	67.30(-16.03)	63.78(-19.55)	67.86(-15.47)	54.27(-29.06)
,	95.23	86.94(-8.29)	89.07(-6.16)	78.15(-17.08)	70.08(-25.15)	93.50(-1.73)	69.87(-25.36)
VB	90.17	69.52(-20.65)	84.21(-5.96)	73.87(-16.30)	76.77(-13.40)	80.49(-9.68)	68.76(-21.41)
VBD	90.97	56.33(-34.64)	85.68(-5.29)	70.51(-20.46)	71.36(-19.61)	70.25(-20.72)	60.42(-30.55)
CC	91.40	68.26(-23.14)	90.09(-1.31)	71.91(-19.49)	68.02(-23.38)	83.39(-8.01)	67.49(-23.91)
VBZ	92.44	74.91(-17.53)	90.40(-2.04)	78.07(-14.37)	80.85(-11.59)	77.62(-14.82)	66.56(-25.88)
VBP	90.02	56.63(-33.39)	84.35(-5.67)	72.21(-17.81)	73.85(-16.17)	71.60(-18.42)	60.45(-29.57)
VBN	88.81	53.78(-35.03)	74.71(-14.10)	67.99(-20.82)	70.24(-18.57)	76.35(-12.46)	55.83(-32.98)
CD	93.10	61.67(-31.43)	74.84(-18.26)	88.25(-4.85)	89.97(-3.13)	89.95(-3.15)	66.87(-26.23)
VBG	92.63	65.44(-27.19)	85.51(-7.12)	78.36(-14.27)	82.53(-10.10)	82.19(-10.44)	66.12(-26.51)
TO	92.38	85.24(-7.14)	87.95(-4.43)	83.07(-9.31)	79.26(-13.12)	78.24(-14.14)	74.63(-17.75)
MD	92.70	65.72(-26.98)	91.93(-0.77)	90.69(-2.01)	90.29(-2.41)	89.55(-3.15)	89.29(-3.41)
PRP\$	92.33	76.90(-15.43)	89.37(-2.96)	88.84(-3.49)	86.96(-5.37)	86.48(-5.85)	80.86(-11.47)
UH	93.66	75.95(-17.71)	87.03(-6.63)	76.42(-17.24)	83.97(-9.69)	85.52(-8.14)	73.79(-19.87)
HYPH	91.98	90.15(-1.83)	92.01(+0.03)	89.90(-2.08)	88.20(-3.78)	92.47(+0.49)	91.23(-0.75)
POS	98.85	93.01(-5.84)	90.80(-8.05)	98.21(-0.64)	97.86(-0.99)	98.37(-0.48)	92.02(-6.83)
'	89.05	76.84(-12.21)	74.25(-14.80)	77.14(-11.91)	71.12(-17.93)	85.48(-3.57)	70.51(-18.54)
"	89.54	76.98(-12.56)	79.12(-10.42)	81.55(-7.99)	68.74(-20.80)	87.44(-2.10)	69.67(-19.87)
WDT	87.26	68.74(-18.52)	85.27(-1.99)	82.54(-4.72)	83.23(-4.03)	80.71(-6.55)	80.34(-6.92)
WP	90.19	76.18(-14.01)	86.99(-3.20)	81.88(-8.31)	83.79(-6.40)	77.35(-12.84)	75.89(-14.30)
WRB	89.15	72.35(-16.80)	87.38(-1.77)	86.68(-2.47)	83.99(-5.16)	79.75(-9.40)	78.44(-10.71)
RP	82.13	74.33(-7.80)	72.20(-9.93)	79.06(-3.07)	82.93(+0.80)	76.13(-6.00)	71.68(-10.45)
:	84.58	66.48(-18.10)	76.22(-8.36)	85.89(+1.31)	80.03(-4.55)	85.33(+0.75)	84.02(-0.56)
JJR	87.18	56.62(-30.56)	71.87(-15.31)	73.65(-13.53)	72.29(-14.89)	72.22(-14.96)	58.62(-28.56)
NNPS	84.39	72.11(-12.28)	75.45(-8.94)	71.02(-13.37)	79.74(-4.65)	79.81(-4.58)	79.96(-4.43)
EX	95.33	89.25(-6.08)	92.18(-3.15)	90.55(-4.78)	90.56(-4.77)	90.77(-4.56)	88.49(-6.84)
JJS	83.41	58.95(-24.46)	77.58(-5.83)	80.64(-2.77)	76.57(-6.84)	79.33(-4.08)	69.87(-13.54)
RBR	87.29	67.71(-19.58)	66.07(-21.22)	66.07(-21.22)	69.96(-17.33)	65.17(-22.12)	58.00(-29.29)
-LRB-	91.54	83.76(-7.78)	89.68(-1.86)	89.51(-2.03)	88.83(-2.71)	88.66(-2.88)	85.96(-5.58)
-RRB-	92.18	83.67(-8.51)	88.78(-3.40)	88.78(-3.40)	87.07(-5.11)	89.97(-2.21)	86.22(-5.96)
\$	97.30	95.38(-1.92)	95.38(-1.92)	94.99(-2.31)	94.80(-2.50)	96.92(-0.38)	94.99(-2.31)
PDT	87.15	84.94(-2.21)	83.54(-3.61)	82.73(-4.42)	81.12(-6.03)	84.74(-2.41)	82.53(-4.62)
RBS	86.31	72.32(-13.99)	72.62(-13.69)	77.68(-8.63)	71.73(-14.58)	77.98(-8.33)	74.41(-11.90)
FW	72.51	64.95(-7.56)	65.98(-6.53)	69.07(-3.44)	76.97(+4.46)	68.04(-4.47)	70.79(-1.72)
NFP	90.56	86.67(-3.89)	88.89(-1.67)	93.33(+2.77)	90.00(-0.56)	92.22(+1.66)	86.67(-3.89)
WP\$	92.99	89.47(-3.52)	95.62(+2.63)	94.74(+1.75)	92.11(-0.88)	94.74(+1.75)	92.98(-0.01)
XX	83.34	81.25(-2.09)	82.29(-1.05)	86.46(+3.12)	83.33(-0.01)	82.29(-1.05)	81.25(-2.09)
SYM	91.11	86.67(-4.44)	88.89(-2.22)	94.45(+3.34)	88.89(-2.22)	93.33(+2.22)	91.11(-0.00)
ADD	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
LS	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
AFX	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)

Table 7: Probing results for POS. STL: the baseline, BERT:pre-trained BERT before fine-tuning.

	STL	BERT	POS	DEP	CON	SRL	MTL-5
GPE	91.40	76.43(-14.97)	82.02(-9.38)	74.70(-16.70)	79.97(-11.43)	84.97(-6.43)	79.73(-11.67)
PERSON	85.90	85.51(-0.39)	85.38(-0.52)	86.67(+0.77)	83.65(-2.25)	83.96(-1.94)	83.60(-2.30)
ORG	73.81	67.80(-6.01)	74.95(+1.14)	69.01(-4.80)	68.04(-5.77)	72.09(-1.72)	69.99(-3.82)
DATE	87.39	63.17(-24.22)	62.51(-24.88)	78.49(-8.90)	73.53(-13.86)	82.25(-5.14)	83.37(-4.02)
CARDINAL	76.33	66.63(-9.70)	66.41(-9.92)	64.81(-11.52)	79.22(+2.89)	73.05(-3.28)	67.06(-9.27)
NORP	93.34	82.28(-11.06)	89.57(-3.77)	88.47(-4.87)	82.72(-10.62)	89.58(-3.76)	90.37(-2.97)
PERCENT	99.81	99.71(-0.10)	98.95(-0.86)	99.71(-0.10)	99.43(-0.38)	99.90(+0.09)	99.71(-0.10)
MONEY	98.30	97.77(-0.53)	96.60(-1.70)	97.98(-0.32)	96.60(-1.70)	97.13(-1.17)	96.07(-2.23)
TIME	79.25	77.36(-1.89)	78.14(-1.11)	72.48(-6.77)	78.77(-0.48)	77.51(-1.74)	77.05(-2.20)
ORDINAL	90.94	89.74(-1.20)	90.26(-0.68)	92.65(+1.71)	90.43(-0.51)	88.03(-2.91)	84.61(-6.33)
LOC	74.86	69.83(-5.03)	76.35(+1.49)	72.81(-2.05)	82.87(+8.01)	74.68(-0.18)	76.91(+2.05)
WORK_OF_ART	78.11	76.51(-1.60)	74.50(-3.61)	76.31(-1.80)	74.50(-3.61)	79.52(+1.41)	67.87(-10.24)
FAC	82.96	77.78(-5.18)	86.42(+3.46)	81.73(-1.23)	74.07(-8.89)	85.43(+2.47)	86.67(+3.71)
QUANTITY	95.55	83.81(-11.74)	95.56(+0.01)	92.70(-2.85)	87.62(-7.93)	94.92(-0.63)	89.84(-5.71)
PRODUCT	85.09	81.58(-3.51)	86.84(+1.75)	88.60(+3.51)	76.75(-8.34)	84.65(-0.44)	78.95(-6.14)
EVENT	77.25	68.25(-9.00)	74.60(-2.65)	71.43(-5.82)	66.14(-11.11)	71.43(-5.82)	70.37(-6.88)
LAW	91.67	92.50(+0.83)	87.50(-4.17)	90.83(-0.84)	86.67(-5.00)	90.00(-1.67)	87.50(-4.17)
LANGUAGE	98.48	95.45(-3.03)	100.00(+1.52)	98.48(-0.00)	98.48(-0.00)	98.48(-0.00)	98.48(-0.00)

Table 8: Probing results for NER. STL: the baseline, BERT:pre-trained BERT before fine-tuning.

	STL	BERT	POS	NER	CON	SRL	MTL-5
prep	72.80	72.62(-0.18)	69.98(-2.82)	70.90(-1.90)	73.11(+0.31)	76.60(+3.80)	78.02(+5.22)
pobj	95.64	96.35(+0.71)	95.01(-0.63)	91.70(-3.94)	94.67(-0.97)	93.45(-2.19)	88.59(-7.05)
nsubj	76.92	76.36(-0.56)	66.45(-10.47)	73.78(-3.14)	73.75(-3.17)	85.61(+8.69)	70.69(-6.23)
det	94.07	93.70(-0.37)	93.22(-0.85)	93.87(-0.20)	95.68(+1.61)	95.72(+1.65)	91.66(-2.41)
root	96.28	96.25(-0.03)	95.83(-0.45)	96.49(+0.21)	95.10(-1.18)	97.77(+1.49)	95.18(-1.10)
nn	89.08	89.47(+0.39)	84.77(-4.31)	86.99(-2.09)	92.59(+3.51)	89.99(+0.91)	89.98(+0.90)
amod	92.76	93.33(+0.57)	91.38(-1.38)	91.63(-1.13)	95.36(+2.60)	92.70(-0.06)	93.50(+0.74)
dobj	93.76	94.55(+0.79)	94.88(+1.12)	93.60(-0.16)	93.46(-0.30)	92.95(-0.81)	89.07(-4.69)
advmod	70.58	69.12(-1.46)	66.11(-4.47)	71.85(+1.27)	69.23(-1.35)	68.96(-1.62)	67.92(-2.66)
aux	84.70	84.53(-0.17)	84.03(-0.67)	85.78(+1.08)	84.22(-0.48)	86.06(+1.36)	84.02(-0.68)
cc	59.17	57.86(-1.31)	58.10(-1.07)	58.73(-0.44)	57.11(-2.06)	54.40(-4.77)	64.18(+5.01)
conj	66.66	64.82(-1.84)	43.20(-23.46)	50.76(-15.90)	65.78(-0.88)	48.50(-18.16)	43.52(-23.14)
dep	41.44	42.77(+1.33)	32.82(-8.62)	41.88(+0.44)	41.22(-0.22)	42.01(+0.57)	38.78(-2.66)
poss	78.20	77.96(-0.24)	79.99(+1.79)	81.39(+3.19)	77.72(-0.48)	72.42(-5.78)	79.02(+0.82)
ccomp	71.24	71.47(+0.23)	50.24(-21.00)	69.80(-1.44)	65.29(-5.95)	66.38(-4.86)	61.40(-9.84)
cop	88.80	87.09(-1.71)	87.47(-1.33)	89.91(+1.11)	88.74(-0.06)	88.04(-0.76)	85.88(-2.92)
mark	89.78	90.88(+1.10)	89.44(-0.34)	91.01(+1.23)	90.29(+0.51)	87.65(-2.13)	89.75(-0.03)
xcomp	68.11	70.53(+2.42)	65.81(-2.30)	72.80(+4.69)	73.20(+5.09)	73.33(+5.22)	71.54(+3.43)
num	84.94	87.61(+2.67)	83.58(-1.36)	87.48(+2.54)	91.50(+6.56)	88.02(+3.08)	90.49(+5.55)
rcmod	51.72	52.62(+0.90)	41.62(-10.10)	41.96(-9.76)	52.22(+0.50)	59.36(+7.64)	51.52(-0.20)
advcl	52.26	54.93(+2.67)	41.78(-10.48)	53.70(+1.44)	49.97(-2.29)	42.06(-10.20)	48.50(-3.76)
neg	82.06	83.40(+1.34)	81.66(-0.40)	82.67(+0.61)	84.82(+2.76)	81.09(-0.97)	78.43(-3.63)
auxpass	97.37	97.51(+0.14)	97.18(-0.19)	97.26(-0.11)	95.74(-1.63)	95.66(-1.71)	96.24(-1.13)
nsubjpass	79.34	83.76(+4.42)	73.26(-6.08)	79.43(+0.09)	73.45(-5.89)	77.24(-2.10)	74.59(-4.75)
possessive	99.23	99.23(-0.00)	99.23(-0.00)	99.20(-0.03)	99.26(+0.03)	99.20(-0.03)	99.33(+0.10)
pcomp	87.55	90.06(+2.51)	83.03(-4.52)	86.95(-0.60)	85.11(-2.44)	84.57(-2.98)	80.15(-7.40)
discourse	74.04	73.80(-0.24)	50.24(-23.80)	62.18(-11.86)	56.72(-17.32)	73.14(-0.90)	52.05(-21.99)
partmod	60.08	59.47(-0.61)	59.52(-0.56)	60.49(+0.41)	63.14(+3.06)	64.85(+4.77)	62.58(+2.50)
appos	54.16	47.88(-6.28)	41.85(-12.31)	47.30(-6.86)	54.65(+0.49)	50.12(-4.04)	43.55(-10.61)
prt	96.44	95.99(-0.45)	96.39(-0.05)	96.69(+0.25)	96.83(+0.39)	95.30(-1.14)	96.09(-0.35)
number	80.75	78.43(-2.32)	77.71(-3.04)	82.86(+2.11)	81.12(+0.37)	81.41(+0.66)	82.57(+1.82)
quantmod	75.49	73.06(-2.43)	72.65(-2.84)	78.32(+2.83)	80.18(+4.69)	73.06(-2.43)	75.41(-0.08)
parataxis	46.48	43.92(-2.56)	30.69(-15.79)	45.24(-1.24)	43.30(-3.18)	47.79(+1.31)	33.69(-12.79)
infmod	71.38	71.20(-0.18)	68.35(-3.03)	70.13(-1.25)	70.49(-0.89)	73.24(+1.86)	68.18(-3.20)
tmod	84.89	87.64(+2.75)	69.23(-15.66)	78.02(-6.87)	76.01(-8.88)	80.31(-4.58)	62.64(-22.25)
expl	85.73	86.27(+0.54)	86.93(+1.20)	86.82(+1.09)	84.75(-0.98)	84.75(-0.98)	84.10(-1.63)
mwe	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
npadvmod	86.24	87.16(+0.92)	88.07(+1.83)	82.42(-3.82)	84.25(-1.99)	83.03(-3.21)	81.19(-5.05)
iobj	93.30	94.02(+0.72)	91.12(-2.18)	89.67(-3.63)	91.85(-1.45)	93.66(+0.36)	88.95(-4.35)
predet	91.72	91.30(-0.42)	91.93(+0.21)	93.17(+1.45)	92.34(+0.62)	87.78(-3.94)	90.47(-1.25)
acomp	89.60	89.17(-0.43)	88.33(-1.27)	89.60(-0.00)	89.81(+0.21)	87.05(-2.55)	89.17(-0.43)
csbj	56.90	57.76(+0.86)	57.19(+0.29)	59.48(+2.58)	53.45(-3.45)	51.72(-5.18)	55.17(-1.73)
preconj	76.81	78.26(+1.45)	65.22(-11.59)	72.46(-4.35)	82.61(+5.80)	83.33(+6.52)	75.36(-1.45)
csbjpass	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	77.78(-22.22)	66.67(-33.33)

Table 9: Probing results for DEP. STL: the baseline, BERT:pre-trained BERT before fine-tuning.

	STL	BERT	POS	NER	DEP	SRL	MTL-5
NP	85.72	55.97(-29.75)	65.53(-20.19)	72.55(-13.17)	67.25(-18.47)	71.80(-13.92)	57.58(-28.14)
ADVP	68.61	68.61(-0.00)	70.25(+1.64)	77.86(+9.25)	70.47(+1.86)	61.15(-7.46)	66.59(-2.02)
ADJP	67.83	53.25(-14.58)	69.33(+1.50)	65.40(-2.43)	57.69(-10.14)	64.42(-3.41)	63.43(-4.40)
VP	79.29	62.83(-16.46)	66.39(-12.90)	70.30(-8.99)	78.07(-1.22)	54.85(-24.44)	69.70(-9.59)
NML	82.92	74.89(-8.03)	81.51(-1.41)	82.17(-0.75)	81.66(-1.26)	84.24(+1.32)	76.02(-6.90)
WHNP	77.87	62.26(-15.61)	75.12(-2.75)	73.27(-4.60)	78.82(+0.95)	70.63(-7.24)	71.14(-6.73)
INTJ	76.61	72.67(-3.94)	84.84(+8.23)	78.11(+1.50)	68.79(-7.82)	73.13(-3.48)	74.11(-2.50)
QP	94.55	90.45(-4.10)	93.49(-1.06)	93.73(-0.82)	91.70(-2.85)	91.65(-2.90)	88.28(-6.27)
WHADVP	78.99	69.43(-9.56)	80.16(+1.17)	75.84(-3.15)	84.04(+5.05)	72.63(-6.36)	76.95(-2.04)
PRT	80.85	79.46(-1.39)	80.90(+0.05)	81.20(+0.35)	83.63(+2.78)	73.71(-7.14)	76.34(-4.51)
PP	73.52	69.86(-3.66)	79.00(+5.48)	78.31(+4.79)	72.83(-0.69)	69.86(-3.66)	71.00(-2.52)
CONJP	90.87	91.78(+0.91)	90.41(-0.46)	85.84(-5.03)	88.58(-2.29)	88.13(-2.74)	87.67(-3.20)
X	59.35	60.98(+1.63)	52.85(-6.50)	53.66(-5.69)	57.73(-1.62)	55.29(-4.06)	56.91(-2.44)
WHADJP	95.37	91.67(-3.70)	93.52(-1.85)	96.29(+0.92)	94.44(-0.93)	94.44(-0.93)	91.67(-3.70)
META	98.61	95.83(-2.78)	97.22(-1.39)	95.83(-2.78)	97.22(-1.39)	97.22(-1.39)	98.61(-0.00)
UCP	100.00	100.00(-0.00)	98.25(-1.75)	100.00(-0.00)	98.25(-1.75)	100.00(-0.00)	100.00(-0.00)
S	96.29	94.44(-1.85)	96.29(-0.00)	94.44(-1.85)	96.29(-0.00)	98.15(+1.86)	94.44(-1.85)
LST	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
FRAG	83.33	100.00(+16.67)	100.00(+16.67)	94.44(+11.11)	94.44(+11.11)	88.89(+5.56)	100.00(+16.67)
SBAR	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
SQ	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
WHPP	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
TOP	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)

Table 10: Probing results for CON. STL: the baseline, BERT:pre-trained BERT before fine-tuning.

	STL	BERT	POS	NER	DEP	CON	MTL-5
ARG1	77.75	71.88(-5.87)	78.78(+1.03)	77.86(+0.11)	77.08(-0.67)	79.58(+1.83)	74.58(-3.17)
ARG0	73.24	74.46(+1.22)	75.95(+2.71)	71.61(-1.63)	71.17(-2.07)	74.16(+0.92)	70.13(-3.11)
ARG2	89.69	80.43(-9.26)	89.81(+0.12)	89.65(-0.04)	90.15(+0.46)	89.86(+0.17)	87.36(-2.33)
ARGM-TMP	73.24	60.55(-12.69)	74.48(+1.24)	70.57(-2.67)	76.58(+3.34)	75.59(+2.35)	68.55(-4.69)
ARGM-DIS	70.40	50.18(-20.22)	69.56(-0.84)	75.67(+5.27)	69.58(-0.82)	77.51(+7.11)	60.41(-9.99)
ARGM-ADV	59.01	40.28(-18.73)	61.31(+2.30)	64.83(+5.82)	56.48(-2.53)	63.88(+4.87)	54.36(-4.65)
ARGM-MOD	82.53	80.46(-2.07)	81.77(-0.76)	84.41(+1.88)	82.29(-0.24)	81.59(-0.94)	78.54(-3.99)
ARGM-LOC	79.45	58.99(-20.46)	78.24(-1.21)	80.68(+1.23)	72.88(-6.57)	79.35(-0.10)	74.22(-5.23)
ARGM-MNR	80.67	69.91(-10.76)	79.33(-1.34)	80.52(-0.15)	80.20(-0.47)	80.89(+0.22)	77.23(-3.44)
ARGM-NEG	86.47	79.39(-7.08)	85.51(-0.96)	90.54(+4.07)	90.81(+4.34)	84.75(-1.72)	88.40(+1.93)
R-ARG1	69.92	67.43(-2.49)	64.77(-5.15)	71.28(+1.36)	74.75(+4.83)	72.09(+2.17)	75.41(+5.49)
R-ARG0	81.20	75.64(-5.56)	76.89(-4.31)	81.71(+0.51)	83.36(+2.16)	80.30(-0.90)	78.36(-2.84)
C-ARG1	51.72	47.46(-4.26)	51.25(-0.47)	46.86(-4.86)	45.51(-6.21)	45.84(-5.88)	46.11(-5.61)
ARGM-PRP	80.48	69.05(-11.43)	78.81(-1.67)	82.06(+1.58)	77.38(-3.10)	83.49(+3.01)	80.16(-0.32)
ARGM-DIR	93.89	89.29(-4.60)	93.73(-0.16)	95.16(+1.27)	92.86(-1.03)	94.92(+1.03)	94.76(+0.87)
ARG3	91.00	82.89(-8.11)	91.98(+0.98)	90.91(-0.09)	87.34(-3.66)	90.91(-0.09)	89.04(-1.96)
ARG4	94.78	93.96(-0.82)	95.51(+0.73)	95.05(+0.27)	94.32(-0.46)	94.32(-0.46)	93.96(-0.82)
ARGM-CAU	62.34	55.37(-6.97)	63.37(+1.03)	66.76(+4.42)	68.74(+6.40)	67.98(+5.64)	65.73(+3.39)
ARGM-PRD	57.72	52.14(-5.58)	61.87(+4.15)	59.66(+1.94)	61.35(+3.63)	59.40(+1.68)	57.85(+0.13)
ARGM-ADJ	85.07	76.00(-9.07)	82.53(-2.54)	82.80(-2.27)	82.67(-2.40)	84.40(-0.67)	82.93(-2.14)
ARGM-EXT	87.37	86.93(-0.44)	90.85(+3.48)	87.80(+0.43)	89.32(+1.95)	88.89(+1.52)	83.44(-3.93)
ARGM-PNC	82.89	69.74(-13.15)	84.21(+1.32)	84.21(+1.32)	82.89(-0.00)	83.77(+0.88)	79.82(-3.07)
ARGM-GOL	86.76	76.71(-10.05)	89.95(+3.19)	88.58(+1.82)	78.54(-8.22)	85.84(-0.92)	84.93(-1.83)
ARGM-LVB	98.59	97.18(-1.41)	98.12(-0.47)	98.12(-0.47)	98.12(-0.47)	97.18(-1.41)	97.18(-1.41)
R-ARGM-LOC	82.05	84.62(+2.57)	78.97(-3.08)	86.16(+4.11)	74.36(-7.69)	73.85(-8.20)	79.49(-2.56)
R-ARGM-TMP	74.08	71.43(-2.65)	69.84(-4.24)	82.01(+7.93)	80.42(+6.34)	75.66(+1.58)	77.25(+3.17)
R-ARG2	73.12	70.97(-2.15)	77.96(+4.84)	77.42(+4.30)	74.19(+1.07)	69.89(-3.23)	74.73(+1.61)
C-ARG2	44.90	36.73(-8.17)	45.58(+0.68)	43.54(-1.36)	43.54(-1.36)	44.22(-0.68)	45.58(+0.68)
C-ARG0	59.80	47.06(-12.74)	61.76(+1.96)	56.86(-2.94)	51.96(-7.84)	53.92(-5.88)	54.90(-4.90)
ARGM-REC	84.31	70.59(-13.72)	85.29(+0.98)	78.43(-5.88)	80.39(-3.92)	85.29(+0.98)	80.39(-3.92)
ARGM-COM	88.89	81.48(-7.41)	90.12(+1.23)	92.59(+3.70)	87.66(-1.23)	87.66(-1.23)	87.66(-1.23)
C-ARGM-ADV	54.55	54.55(-0.00)	48.48(-6.07)	54.55(-0.00)	54.55(-0.00)	51.52(-3.03)	57.58(+3.03)
R-ARGM-MNR	72.73	63.64(-9.09)	63.64(-9.09)	81.82(+9.09)	69.70(-3.03)	78.79(+6.06)	66.67(-6.06)
ARG5	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-TMP	71.43	42.86(-28.57)	61.90(-9.53)	71.43(-0.00)	57.14(-14.29)	61.90(-9.53)	66.67(-4.76)
R-ARGM-CAU	91.67	75.00(-16.67)	83.33(-8.34)	100.00(+8.33)	100.00(+8.33)	91.67(-0.00)	91.67(-0.00)
C-ARGM-CAU	66.67	66.67(-0.00)	66.67(-0.00)	77.78(+11.11)	66.67(-0.00)	66.67(-0.00)	77.78(+11.11)
C-ARG3	100.00	100.00(-0.00)	88.89(-11.11)	88.89(-11.11)	88.89(-11.11)	88.89(-11.11)	100.00(-0.00)
R-ARG4	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-EXT	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
R-ARGM-ADV	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-LOC	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-MNR	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-MOD	100.00	50.00(-50.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
R-ARGM-PRP	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
ARGM-PRX	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
ARGA	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
R-ARGM-DIR	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
R-ARGM-PRD	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-PRD	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
R-ARGM-EXT	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
C-ARGM-PRP	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
R-ARG3	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
ARGM-PRR	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)
ARGM-DSP	100.00	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)	100.00(-0.00)

Table 11: Probing results for SRL. STL: the baseline, BERT:pre-trained BERT before fine-tuning.

A.4 Results for Other Transformers

We also applied our pruning and probing methods on 3 recent transformer encoders, RoBERTa (Liu et al., 2019c), ELECTRA (Clark et al., 2020) and DeBERTa (He et al., 2020), to further demonstrate the generality of our hypothesis. For all of them, we use the base size version which has 12 layers with 144 attention heads in total. Although we did not tune hyper-parameters specifically for any of them and re-used the same hyper-parameters of BERT, their results turned out to be as interesting as BERT results.

Their STL and MTL results are shown in Table 12. Unsurprisingly, MTL-5 is outperformed by its single-task counterparts for all tasks and for all transformer encoders, raising the dilemma behind transformer-based MTL.

Their pruning results are shown in Table 13. Although the results could be better tuned for each transformer encoder, our DP strategy is still able to prune roughly 50% heads while keeping comparable performance.

Their visualization of head utilization is illustrated in Figure 9, 11 and 13. Similar patterns among each transformer encoder can be observed, supporting our claim that the MTL-DP model re-uses a very similar set of heads used by the STL-DP models.

Their probing results are illustrated in Figure 10, 12 and 14, which also align with our findings. Specifically, the DEP probing results on transformer encoders are already very high even without fine-tuning on actual dependency treebanks (Figure 10f, 12f and 14f), demonstrating the pluripotency of *stem cells*. Other figures show similar patterns that *stem cell* heads get specialized during STL but lose specialities in MTL, supporting our *Stem Cell Hypothesis*.

	POS	NER	DEP	CON	SRL	MTL-5
POS	98.33 ± 0.01	98.30 ± 0.01	98.34 ± 0.00	98.33 ± 0.02	98.31 ± 0.02	98.29 ± 0.02
NER	89.44 ± 0.04	88.93 ± 0.16	89.84 ± 0.14	89.65 ± 0.02	89.37 ± 0.19	89.66 ± 0.17
DEP	94.48 ± 0.05	94.46 ± 0.04	94.56 ± 0.05	94.55 ± 0.02	94.44 ± 0.02	94.38 ± 0.03
CON	94.82 ± 0.01	94.76 ± 0.04	94.88 ± 0.03	94.89 ± 0.05	94.88 ± 0.01	94.75 ± 0.01
SRL	84.04 ± 0.05	83.31 ± 0.11	84.13 ± 0.08	84.14 ± 0.02	83.52 ± 0.07	83.56 ± 0.10

(a) RoBERTa (Liu et al., 2019c) performance.

	POS	NER	DEP	CON	SRL	MTL-5
POS	97.94 ± 0.04	97.97 ± 0.02	97.94 ± 0.01	97.95 ± 0.01	97.96 ± 0.05	97.93 ± 0.01
NER	88.62 ± 0.03	88.22 ± 0.05	88.46 ± 0.20	88.82 ± 0.14	88.25 ± 0.11	88.38 ± 0.04
DEP	94.52 ± 0.08	94.47 ± 0.08	94.61 ± 0.01	94.65 ± 0.03	94.53 ± 0.07	94.39 ± 0.04
CON	95.01 ± 0.01	95.04 ± 0.02	95.10 ± 0.06	95.14 ± 0.03	95.08 ± 0.07	94.91 ± 0.04
SRL	84.35 ± 0.07	83.44 ± 0.05	84.55 ± 0.06	84.61 ± 0.11	83.81 ± 0.11	84.09 ± 0.06

(b) ELECTRA (Clark et al., 2020) performance.

	POS	NER	DEP	CON	SRL	MTL-5
POS	98.33 ± 0.02	98.26 ± 0.02	98.32 ± 0.01	98.32 ± 0.01	98.31 ± 0.02	98.30 ± 0.02
NER	89.29 ± 0.18	89.02 ± 0.21	89.67 ± 0.09	89.65 ± 0.25	89.29 ± 0.36	89.28 ± 0.18
DEP	94.50 ± 0.04	94.54 ± 0.02	94.64 ± 0.03	94.61 ± 0.07	94.56 ± 0.04	94.42 ± 0.03
CON	94.91 ± 0.04	94.94 ± 0.02	94.96 ± 0.01	95.06 ± 0.01	94.98 ± 0.03	94.82 ± 0.03
SRL	84.26 ± 0.05	83.42 ± 0.15	84.34 ± 0.04	84.36 ± 0.09	83.62 ± 0.06	83.37 ± 0.01

(c) DeBERTa (He et al., 2020) performance.

Table 12: Performance of single-task learning (main diagonal highlighted in gray), multi-task learning on all 5 tasks (MTL-5), and multi-task learning on every pair of the tasks (non-diagonal cells; e.g., DEP'th row in NER'th column is the DEP result of the joint model between DEP and NER).

	Performance				% of Attention Heads Kept	
	STL	STL-DP	MTL	MTL-DP	STL-DP	MTL-DP
POS	98.33 ± 0.01	98.35 ± 0.02	98.29 ± 0.02	98.38 ± 0.01	41.90 ± 0.40	53.47 ± 1.39
NER	88.93 ± 0.16	89.08 ± 0.16	89.66 ± 0.17	89.52 ± 0.05	58.33 ± 3.48	53.47 ± 1.39
DEP	94.56 ± 0.05	94.29 ± 0.01	94.38 ± 0.03	94.48 ± 0.03	64.58 ± 1.21	53.47 ± 1.39
CON	94.89 ± 0.05	94.59 ± 0.05	94.75 ± 0.01	94.78 ± 0.06	55.33 ± 2.00	53.47 ± 1.39
SRL	83.52 ± 0.07	83.53 ± 0.06	83.56 ± 0.10	83.57 ± 0.14	45.14 ± 0.70	53.47 ± 1.39

(a) RoBERTa (Liu et al., 2019c) pruning results.

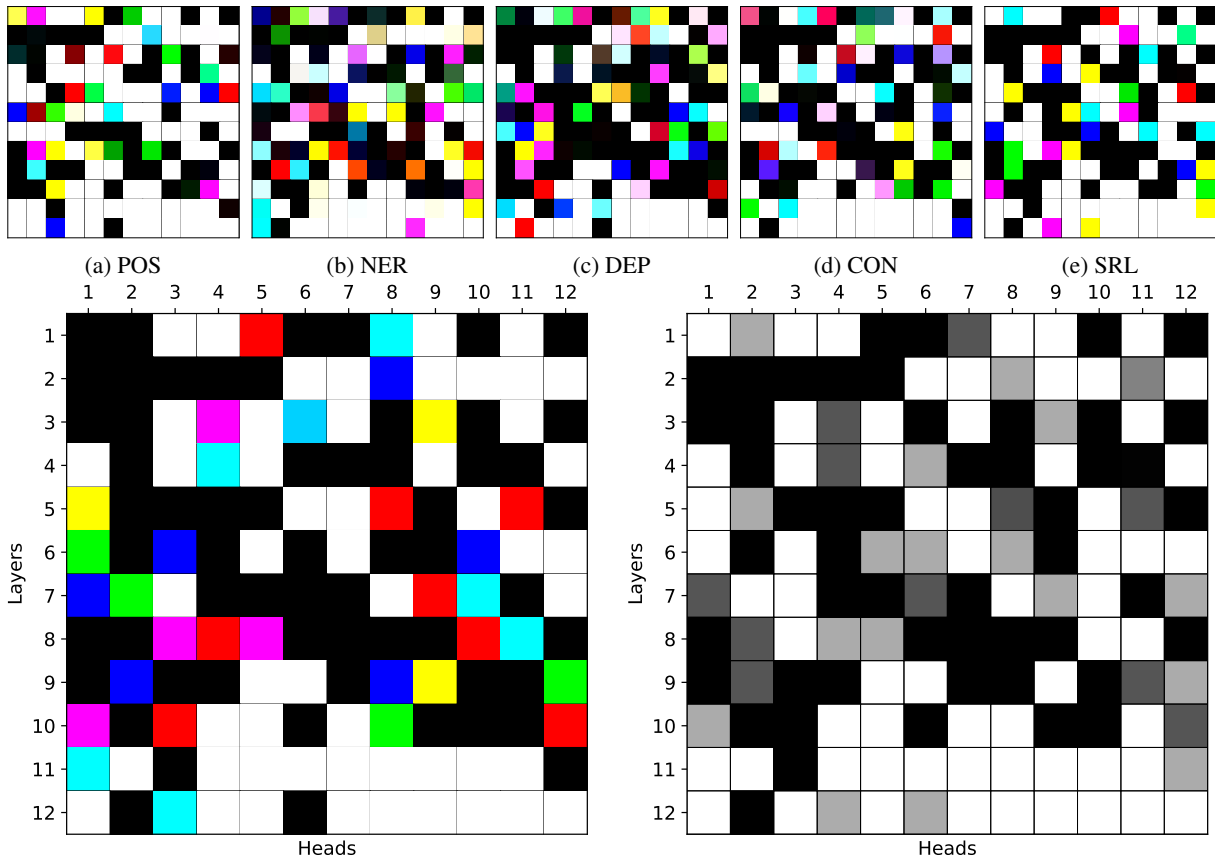
	Performance				% of Attention Heads Kept	
	STL	STL-DP	MTL	MTL-DP	STL-DP	MTL-DP
POS	97.94 ± 0.04	98.01 ± 0.01	97.93 ± 0.01	98.03 ± 0.02	35.88 ± 2.12	44.67 ± 2.44
NER	88.22 ± 0.05	88.04 ± 0.24	88.38 ± 0.04	88.11 ± 0.22	58.33 ± 2.78	44.67 ± 2.44
DEP	94.61 ± 0.01	94.36 ± 0.04	94.39 ± 0.04	94.40 ± 0.01	58.80 ± 1.06	44.67 ± 2.44
CON	95.14 ± 0.03	94.82 ± 0.02	94.91 ± 0.04	94.86 ± 0.07	52.78 ± 2.50	44.67 ± 2.44
SRL	83.81 ± 0.11	83.71 ± 0.10	84.09 ± 0.06	84.07 ± 0.06	41.20 ± 0.40	44.67 ± 2.44

(b) ELECTRA (Clark et al., 2020) pruning results.

	Performance				% of Attention Heads Kept	
	STL	STL-DP	MTL	MTL-DP	STL-DP	MTL-DP
POS	98.33 ± 0.02	98.37 ± 0.01	98.30 ± 0.02	98.39 ± 0.01	50.93 ± 3.13	66.67 ± 1.84
NER	89.02 ± 0.21	88.72 ± 0.18	89.28 ± 0.18	89.08 ± 0.07	66.67 ± 3.67	66.67 ± 1.84
DEP	94.64 ± 0.03	94.42 ± 0.02	94.42 ± 0.03	94.60 ± 0.01	54.63 ± 2.00	66.67 ± 1.84
CON	95.06 ± 0.01	94.93 ± 0.03	94.82 ± 0.03	94.94 ± 0.13	61.11 ± 1.20	66.67 ± 1.84
SRL	83.62 ± 0.06	83.78 ± 0.13	83.37 ± 0.01	83.37 ± 0.00	55.09 ± 0.40	66.67 ± 1.84

(c) DeBERTa (He et al., 2020) pruning results.

Table 13: Results of single-task learning (STL), STL with static pruning (STL-SP) and multi-task learning on the 5 tasks with/without dynamic pruning (MTL/MTL-DP). The STL Performance column is equivalent to the main diagonal in Table 12.



(f) 3-run utilization of the MTL-DP model, where each run is encoded in a RGB channel. Darker indicates higher utilization.

(g) Average head utilization rates among the 5 tasks in 3 runs. Darker cells indicate higher utilization rates.

Figure 9: Head utilization of the RoBERTa (Liu et al., 2019c) STL-DP models (a - e, g) and the MTL-DP model (f).

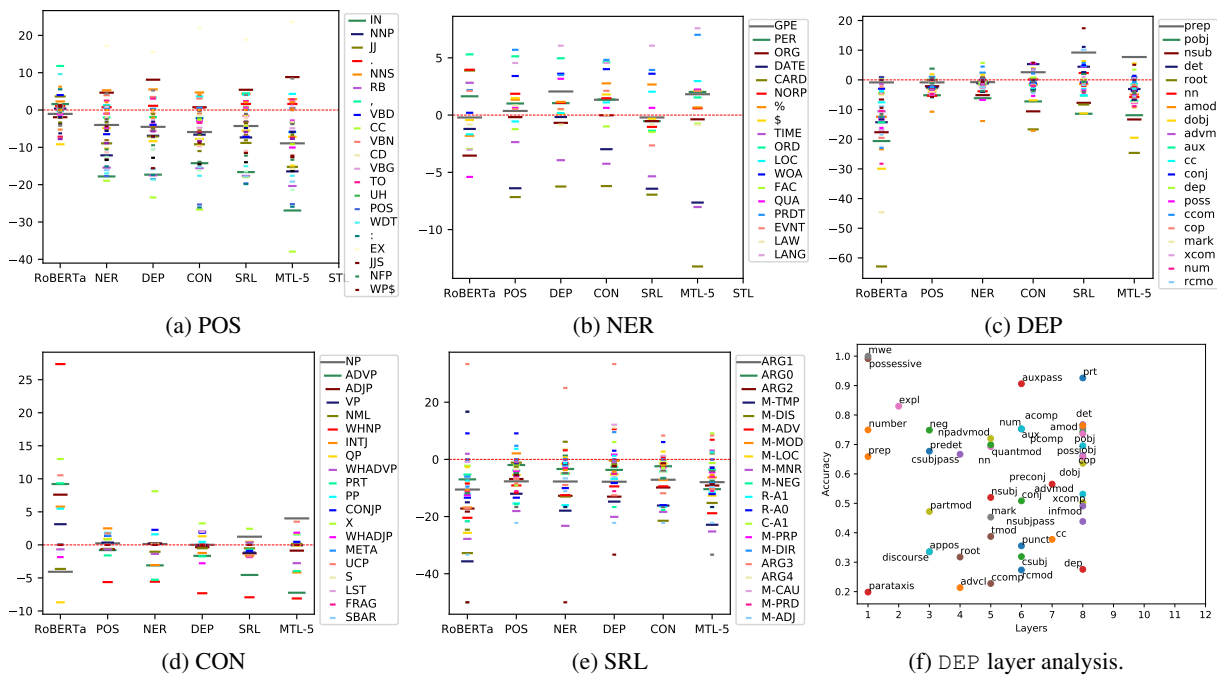





Figure 10: The RoBERTa (Liu et al., 2019c) probing results comparison (a - e) and layer analysis of pre-trained heads (g).



(f) 3-run utilization of the MTL-DP model, where each run is encoded in a RGB channel. Darker indicates higher utilization.



(g) Average head utilization rates among the 5 tasks in 3 runs. Darker cells indicate higher utilization rates.

Figure 11: Head utilization of the ELECTRA (Clark et al., 2020) STL-DP models (a - e, g) and the MTL-DP model (f).

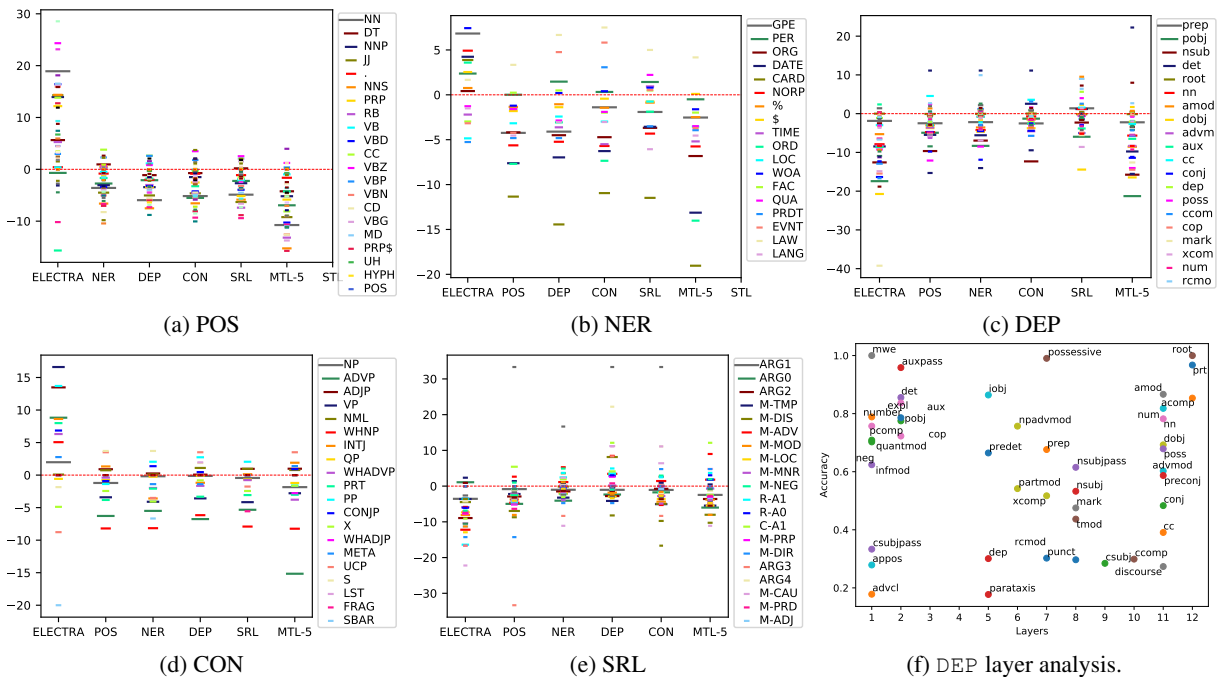





Figure 12: The ELECTRA (Clark et al., 2020) probing results comparison (a - e) and layer analysis of pre-trained heads (g).



(f) 3-run utilization of the MTL-DP model, where each run is encoded in a RGB channel. Darker indicates higher utilization.

(g) Average head utilization rates among the 5 tasks in 3 runs. Darker cells indicate higher utilization rates.

Figure 13: Head utilization of the DeBERTa (He et al., 2020) STL-DP models (a - e, g) and the MTL-DP model (f).

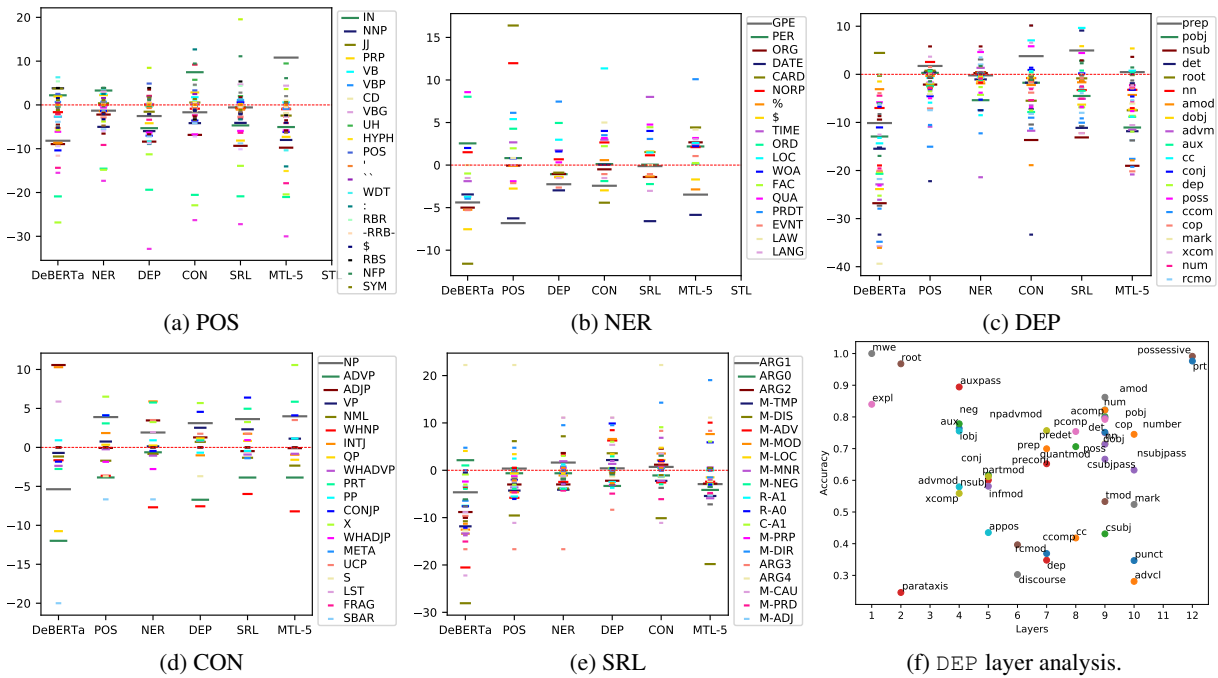


Figure 14: The DeBERTa (He et al., 2020) probing results comparison (a - e) and layer analysis of pre-trained heads (g).