

Discovering Salient Neurons in deep NLP models

Nadir Durrani

*Qatar Computing Research Institute
Hamad Bin Khalifa University
Doha, Qatar*

NDURRANI@HBKU.EDU.QA

Fahim Dalvi

*Qatar Computing Research Institute
Hamad Bin Khalifa University
Doha, Qatar*

FAIMADUDDIN@HBKU.EDU.QA

Hassan Sajjad

*Faculty of Computer Science
Dalhousie University
Halifax, Nova Scotia, Canada*

HSAJJAD@DAL.CA

Editor: Alexander Clark

Abstract

While a lot of work has been done in understanding representations learned within deep NLP models and what knowledge they capture, work done towards analyzing individual neurons is relatively sparse. We present a technique called **Linguistic Correlation Analysis** to extract salient neurons in the model, with respect to any extrinsic property, with the goal of understanding how such knowledge is preserved within neurons. We carry out a fine-grained analysis to answer the following questions: (i) can we identify subsets of neurons in the network that learn a specific linguistic property? (ii) is a certain linguistic phenomenon in a given model localized (encoded in few individual neurons) or distributed across many neurons? (iii) how redundantly is the information preserved? (iv) how does fine-tuning pre-trained models towards downstream NLP tasks impact the learned linguistic knowledge? (v) how do models vary in learning different linguistic properties? Our data-driven, quantitative analysis illuminates interesting findings: (i) we found small subsets of neurons that can predict different linguistic tasks; (ii) neurons capturing basic lexical information, such as suffixation, are localized in the lowermost layers; (iii) neurons learning complex concepts, such as syntactic role, are predominantly found in middle and higher layers; (iv) salient linguistic neurons are relocated from higher to lower layers during transfer learning, as the network preserves the higher layers for task-specific information; (v) we found interesting differences across pre-trained models regarding how linguistic information is preserved within them; and (vi) we found that concepts exhibit similar neuron distribution across different languages in the multilingual transformer models. Our code is publicly available as part of the NeuroX toolkit (Dalvi et al., 2023).¹

Keywords: Neuron Analysis, Representation Analysis, Interpretability, Explainable AI

1. <https://github.com/fdalvi/NeuroX>

1. Introduction

The advent of deep neural networks and their opaque nature has prompted a new field of research focused on analyzing these models. However, this revolution comes at the expense of transparency. Unlike traditional models that rely on manually designed features, we now face a lack of understanding regarding two key aspects: i) what knowledge the underlying representations capture, and ii) how this information is utilized in the decision-making process of deep NLP models. Addressing these questions is crucial for understanding the inner workings of neural networks and ensuring fairness, trust, and accountability when applying AI solutions. To gain insights into these questions, researchers have pursued two main avenues of investigation: i) Concept Analysis and ii) Attribution Analysis. The former thrives on post-hoc decomposability, analyzing representations to uncover linguistic (and non-linguistic) phenomena learned as the network is trained for any NLP task (Belinkov et al., 2017a; Conneau et al., 2018; Tenney et al., 2019). The latter characterizes the role of model components in a specific prediction (Linzen et al., 2016b; Gulordava et al., 2018; Marvin and Linzen, 2018). Recently, efforts have been made to connect these two lines of work and determine whether the encoded knowledge is actually used by the model (Feder et al., 2021; Elazar et al., 2021).

Concept Analysis aims at understanding the learned representations with respect to human interpretable concepts. For example analyzing if the model is aware of language morphology, anaphora or more subtle concepts like gender etc. Such an analysis can be potentially useful for not only debugging the systems for errors (Ma et al., 2018; Lertvittayakumjorn et al., 2020), but improving the design choices and system performance (Wambsganss et al., 2021). To this end the researchers have analyzed contextualized representations trained within deep NLP models w.r.t concepts varying from word morphology (Vylomova et al., 2016; Dalvi et al., 2017; Durrani et al., 2019) to high level concepts such as structure (Shi et al., 2016; Linzen et al., 2016a) and semantics (Qian et al., 2016b; Belinkov et al., 2017b) or more generic properties such as sentence length (Adi et al., 2016). Compared to the holistic view that we obtain from representation analysis, neuron analysis offers several benefits. By analyzing individual neurons, we gain a more fine-grained understanding of the network’s representations. This allows us to identify specific neurons or groups of neurons that capture particular phenomena or concepts (Karpathy et al., 2015; Li et al., 2016; Kádár et al., 2017; Lakretz et al., 2019). Such insights go beyond understanding the mechanics of neural networks and have numerous applications. For example understanding which neurons are responsible for certain behaviors enables us to intervene and manipulate the network’s output accordingly (Bau et al., 2019). Neuron analysis can facilitate model distillation by identifying less salient neurons (Rethmeier et al., 2020). It can contribute to efficient feature-based transfer learning (Dalvi et al., 2020) by understanding the functional role of different neurons. Finally, how information is distributed across the architectures can guide architecture search.

We present a method called as the *Linguistic Correlation Analysis (LCA)*, to extract salient neurons in the model with respect to any extrinsic property. Our methodology is a 3-step process — given a neural network model: i) use the pre-trained model to extract contextualized feature representations for words, ii) train a classifier using the extracted features to predict a property of interest, iii) extract a ranking based on the weights as-

signed to each neuron. The ranking serves as a proxy to measure the importance of neurons with respect to the understudied property. We select top $N\%$ neurons from the ranked list, that maintain oracle performance i.e. accuracy on the task using the entire representation.

The LCA method is based on the diagnostic classifier framework (Hupkes et al., 2017; Belinkov et al., 2017a). However, rather than examining the overall representations learned by the neural network, we use it to discover individual neurons w.r.t their role in encoding specific linguistic features.

We apply the LCA method to delve into the core linguistic phenomena encompassing morphology, syntax, and semantics, essential aspects of language learning, deemed crucial for any sophisticated NLP task. However, our method can be seamlessly used for any task-specific concept where annotated data is available. We analyze the representations trained from neural language models of which we study the recent transformer based models. We use LCA to understand: i) how concepts are learned within neurons and ii) how these neurons distribute across the network. More specifically we probe for the following questions:

- Question: Do the individual neurons in the transformers capture linguistic information and which parts of the network learn more about certain linguistic phenomena?
- Finding: Neurons that capture tasks such as predicting shallow lexical information (such as suffixes) or word morphology are predominantly found in the lower layers and those learning more complex properties such as syntactic information are found in the higher layers. [Section 5.1]
- Question: Is certain linguistic phenomenon in a given model localized or distributed across many neurons? And how redundantly is the information preserved?
- Finding: Closed-class properties like *interjections* (morphological category) and *Months of year* (semantic category) are concentrated within a smaller number of neurons, whereas open-class properties such as nouns (morphological category) and *location* (semantic category) are spread across multiple neurons. [Section 5.2]
- Finding: A subset of neurons that learn a particular concept can be discovered, however, concepts are redundantly learned and many subsets of neurons that carry the same information coexist in the network. [Sections 4.1,5.3]
- Question: How do various architectures differ in learning these properties?
- Finding: Some models (such as BERT) are more distributed while others (such as XLNet) exhibit disjoint set of neurons designated for different linguistic properties. [Section 5.1]
- Question: How does the information redistributes across the layers during transfer learning?
- Finding: The differences between architectures became more pronounced when the models were tuned towards the GLUE tasks. We observed that linguistic knowledge regresses to lower layers in RoBERTa and XLNet, unlike BERT, where it is still preserved in the higher layers. [Section 5.4]

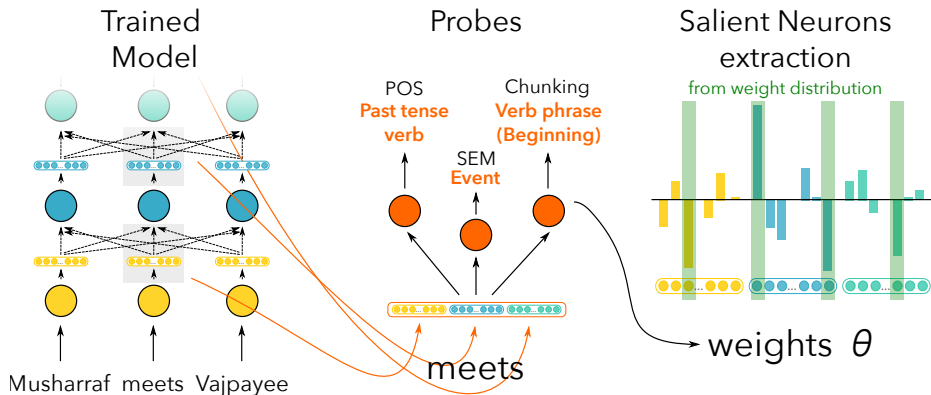


Figure 1: Linguistic Correlation Analysis: Extract neuron activations from a trained model, train a classifier and use weights of the classifier to extract salient neurons.

- Question: Do multilingual models with shared knowledge exhibit distinct behavior?
- Finding: Multilingual models exhibit a similarity in distributing neurons across layers, when learning concepts in different languages. [Section 5.5]

This paper builds upon the work presented in Dalvi et al. (2019) by expanding on several aspects: i) we extend the analysis towards transformer models including multilingual BERT and XLM-RoBERTa (Conneau et al., 2020), ii) we used the algorithm to discover minimal set of salient neurons that learn a linguistic property and carry out several analyses on top of this. For example: a) to understand how salient neurons distribute across the layers and linguistic properties, b) how redundantly is the information preserved, c) how it redistributes when the models are fine-tuned towards a downstream task.

2. Linguistic Correlation Analysis

The LCA algorithm follows a three step process (See Figure 1): i) first we extract activations from the trained neural network model, ii) we then use these as static features to train an auxiliary classifier towards an extrinsic property of interest, iii) lastly we use the weights of the trained classifier as a proxy to the describe saliency of a neuron w.r.t the property.

2.1 Model

Formally, consider a pre-trained neural language model \mathbf{M} with L layers: $\{l_1, l_2, \dots, l_L\}$. Given a dataset $\mathbb{W} = \{w_1, w_2, \dots, w_N\}$ with the corresponding annotations $\mathbb{T} = \{t_{w_1}, t_{w_2}, \dots, t_{w_N}\}$, we map each word w_i in the data \mathbb{W} to a sequence of latent representations: $\mathbb{W} \xrightarrow{\mathbf{M}} \mathbf{z} = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$. The representations can be extracted from the entire model, from an individual layer or any component of the model (attention head, FFN layer etc). The probe is trained by minimizing the following loss function:

$$\mathcal{L}(\theta) = - \sum_i \log P_{\theta}(t_{w_i} | w_i) + \lambda_1 \|\theta\|_1 + \lambda_2 \|\theta\|_2^2$$

Algorithm 1 Probe training (Section 2.1)

```

1: function TRAINPROBE( $X, y, \lambda_1, \lambda_2$ )
2:   Initialize learning rate  $\eta = 0.001$ , number of epochs  $N = 10$ 
3:   Initialize parameters  $\theta$  (weights  $W$  and bias  $b$ ) randomly
4:   for epoch = 1 to  $N$  do
5:     for each training example  $(w_i, t_{w_i})$  in  $X$  do  $\triangleright$  Can be done over mini-batches
        as well
6:        $L(\theta) = -\sum_i \log P_\theta(t_{w_i}|w_i) + \lambda_1 \|\theta\|_1 + \lambda_2 \|\theta\|_2^2$   $\triangleright$  Equation from 2.1
7:        $d\theta \leftarrow \frac{\partial L(\theta)}{\partial \theta}$   $\triangleright$  Compute the gradient of the loss function with respect to  $W$ 
        and  $b$ 
8:        $\theta \leftarrow \theta - \eta \cdot d\theta$   $\triangleright$  Update the parameters  $W$  and  $b$ 
9:     end for
10:  end for
11:  return probe $_\theta$   $\triangleright$  trained probe
12: end function

```

where $P_\theta(t_{w_i}|w_i) = \frac{\exp(\theta_i \cdot \mathbf{z}_i)}{\sum_{t'} \exp(\theta_{t'} \cdot \mathbf{z}_i)}$ is the probability that word i is assigned property t_{w_i} . The weights $\theta \in \mathbb{R}^{D \times |T|}$ are learned with gradient descent (Algorithm 1). Here D is the dimensionality of the latent representations \mathbf{z}_i and T is the set of tags (properties) in the linguistic tag set, which the classifier is predicting, $|T|$ is the number of such tags. The trained probe can also be represented as a function $X \mapsto y$ where X is the combination of all the word representations (w_i) and y is the list of associated labels from T . The terms $\lambda_1 \|\theta\|_1$ and $\lambda_2 \|\theta\|_2^2$ correspond to $L1$ and $L2$ regularization. This combination also known as elastic-net (Zou and Hastie, 2005). It is useful to strike a balance between identifying very focused localized features ($L1$) versus distributed neurons ($L2$). We use a **grid search** algorithm to find the most appropriate set of lambda values.

2.2 Grid Search

To find the best values for λ_1 and λ_2 , we conduct a weight ablation experiment on the trained classifier using a search algorithm (refer to Algorithm 2). Following classifier training, we identify the top and bottom M features from a ranked list generated by a neuron ranking algorithm. The remaining features are then set to zero. The score for each pair of (λ_1, λ_2) is calculated as follows:

$$S(\lambda_1, \lambda_2) = \alpha(A_t - A_b) - \beta(A_z - A_l)$$

here, A_t represents the accuracy of the classifier retaining top neurons and masking the rest, A_b is the accuracy retaining bottom neurons, A_z stands for the accuracy of the classifier trained using all neurons but without regularization, and A_l represents the accuracy with the current lambda set. The first term ensures the selection of a lambda set where accuracies

Algorithm 2 Grid Search (Section 2.2)

```

1: function GRIDSEARCH( $X, y, \alpha, \beta, M$ )
2:   searchList  $\leftarrow [1^{-6}, 1^{-5}, 1^{-4}, 1^{-3}, 1^{-2}, 1^{-1}, 0]$ 
3:    $\lambda_1^*, \lambda_2^* \leftarrow \text{null}$  ▷ Best regularization parameters
4:   bestScore  $\leftarrow -\infty$ 
5:   for  $\lambda_1$  in searchList do
6:     for  $\lambda_2$  in searchList do
7:       probe  $\leftarrow \text{TRAINPROBE}(X, y, \lambda_1, \lambda_2)$ 
8:       rankedNeurons  $\leftarrow \text{GETNEURONRANKING}(\text{probe})$ 
9:        $X_{\text{top}} \leftarrow \text{ABLATEBOTTOMNEURONS}(X, \text{rankedNeurons}, M)$  ▷ Set all
       neurons to zero except top  $M$  neurons
10:       $X_{\text{bottom}} \leftarrow \text{ABLATETOPNEURONS}(X, \text{rankedNeurons}, M)$  ▷ Set all neurons
       to zero except bottom  $M$  neurons
11:      probetop  $\leftarrow \text{TRAINPROBE}(X_{\text{top}}, y, \lambda_1, \lambda_2)$ 
12:      probebottom  $\leftarrow \text{TRAINPROBE}(X_{\text{bottom}}, y, \lambda_1, \lambda_2)$ 
13:      probenoreg  $\leftarrow \text{TRAINPROBE}(X, y, 0, 0)$ 
14:      score  $\leftarrow \alpha \cdot (\text{A}(\text{probe}_{\text{top}}) - \text{A}(\text{probe}_{\text{bottom}})) - \beta \cdot (\text{A}(\text{probe}_{\text{noreg}}) - \text{A}(\text{probe}))$ 
       ▷ A computes the accuracy of the given probe
15:      if score > bestScore then
16:        bestScore  $\leftarrow$  score
17:         $\lambda_1^* \leftarrow \lambda_1$ 
18:         $\lambda_2^* \leftarrow \lambda_2$ 
19:      end if
20:    end for
21:  end for
22:  return  $\lambda_1^*, \lambda_2^*$ 
23: end function

```

of top and bottom neurons are significantly different, while the second term ensures a preference for weights that result in minimal loss in classifier accuracy due to regularization.²

2.3 Neuron Ranking Algorithm

Deriving neuron rankings is analogous to subset selection, which is a well-studied field in machine learning. Classical subset selection methods either enumerate all possible subsets to select the best features, or perform step wise selection of features by training a model to find the next best feature at each step (Hastie et al., 2013, Section 3.3). However, these techniques become infeasible when the number of total features is large, as in the case of neuron selection. In this work, we derive inspiration of Shrinkage Methods (Hastie et al., 2013, Section 3.4) and use the internal weights of a trained classifier to extract top neurons (and subsequently a ranking of neurons). This avoids any expensive computation

2. Higher L1 values encourage sparsity in the classifier, leading to the selection of a smaller set of focused neurons. However, this often results in poorer task performance. We set $M = 20\%$ and $\alpha, \beta = 0.5$.

Algorithm 3 Neuron Ranking Extraction (Section 2.3)

```

1: function GETNEURONRANKING(probe)
2:   ordering  $\leftarrow$  []  $\triangleright$  ordering will store the neurons in order of decreasing importance
3:   for  $p = 1$  to 100 by  $\alpha$  do  $\triangleright p$  is the percentage of the weight mass. We start with a very
      small value and incrementally move towards 100%.
4:     tnpt  $\leftarrow$  GETTOPNEURONSPERTAG( $\theta, p$ )  $\triangleright$  tnpt contains the top neurons per tag
      using the threshold  $p$ 
5:     topNeurons  $\leftarrow \bigcup_{i=1}^L$  tnpt $i$ 
6:     newNeurons  $\leftarrow$  topNeurons  $\setminus$  ordering
7:     ordering.append(newNeurons)
8:   end for
9:   return ordering
10: end function

```

per feature, and also allows the usage of a single trained model to extract independent rankings for each class in a multi-class scenario.

Formally, given the trained weights of the classifier denoted as $\theta \in \mathbb{R}^{D \times |T|}$, our objective is to establish a ranking for the D neurons within the model \mathbb{M} . For a specific property $\mathbf{t} \in T$, we sort the weights $\theta_{\mathbf{t}} \in \mathbb{R}^D$ based on their absolute values, sorting them in descending order. Consequently, the neuron with the highest absolute weight in $\theta_{\mathbf{t}}$ is positioned at the top of our ranking. We identify the top n neurons (pertaining to the specific property in consideration) that cumulatively contribute to some percentage of the total weight mass as *salient neurons*. To extract a ranking of neurons w.r.t. all of the tags in T , we use an iterative process described in Algorithm 3. We start with a small percentage of the total weight mass and choose the most salient neurons for each tag \mathbf{t} , and increase this percentage iteratively, adding newly discovered top neurons to our ordering. Hence, the salient neurons for each label \mathbf{t} appears at the top of the ordering. The order in which the neurons are discovered indicates their importance to the property set T .

2.4 Minimal Neuron Selection

Our approach involves an iterative process to identify minimal neurons for any downstream task: i) train a classifier to predict the task using all the neurons (call it *Oracle*), ii) obtain a neuron ranking based on the ranking algorithm described above, iii) select the top N neurons from the ranked list and retrain a classifier using these selected neurons. This process is repeated, with N gradually increasing by 1% at each iteration, until the classifier achieves an accuracy close to, but not less than a specified threshold compared to the performance of the Oracle. In our experiments, we utilized a threshold of 1% loss compared to Oracle, to guide this iterative process.

2.5 Selectivity

A potential pitfall in the approach of probing classifiers is the challenge of determining whether the probe faithfully reflects the underlying property of the representation or simply

Algorithm 4 Minimal Neuron Subset (Section 2.4)

```

1: function GETMINIMUMNEURONS( $X, y, \delta$ )
2:    $\lambda_1^*, \lambda_2^* \leftarrow$  GRIDSEARCH( $X, y, 0.5, 0.5, 20\%$ )
3:   probe  $\leftarrow$  TRAINPROBE( $X, y, \lambda_1^*, \lambda_2^*$ )
4:    $O_{acc} \leftarrow$  A(probe) ▷ Oracle accuracy
5:   rankedNeurons  $\leftarrow$  GETNEURONRANKING(probe)
6:    $N \leftarrow 0$  ▷ Number of selected neurons
7:   repeat
8:      $N \leftarrow N + 1$ 
9:      $X_{top} \leftarrow$  SELECTTOPNEURONS( $X, rankedNeurons, N$ ) ▷ Select top  $N$  neurons
10:    probe_selected  $\leftarrow$  TRAINPROBE( $X_{top}, y, \lambda_1^*, \lambda_2^*$ )
11:     $S_{acc} \leftarrow$  A(probe_selected) ▷ Selected neurons accuracy
12:    until  $N < \text{length}(\text{rankedNeurons})$  or  $O_{acc} - S_{acc} < \delta$ 
13:    return  $N$ 
14: end function

```

learns the specific task. Hewitt and Liang (2019) introduced control tasks to investigate the influence of training data and lexical memorization on probing experiments. They propose *Selectivity* as a criterion to put a “linguistic task’s accuracy in context with the probe’s capacity to memorize from word types”. An effective probe should exhibit high accuracy in the linguistic task and low accuracy in the control task. We used selectivity to ensure that our probes are a true reflection of the learned representation.

The control tasks for our probing classifiers are defined by mapping each word type x_i to a randomly sampled behavior $C(x_i)$, from a set of numbers $\{1 \dots T\}$ where T is the size of tag set to be predicted in the linguistic task. The sampling is done using the empirical token distribution of the linguistic task, so the marginal probability of each label is similar. The selectivity is defined as: $\delta = A_t - A_{t_c}$ i.e. difference of classifier accuracy when trained towards the task at hand and accuracy when trained towards control task drawn from the same distribution.

3. Experimental Setup

Pre-trained Neural Language Models: We experimented with 3 transformer models: BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019b) and XLNet (Yang et al., 2019) using the base versions (13 layers and 768 dimensions). This architectural choice allowed for a compelling comparison between auto-encoder and auto-regressive models. Additionally, we explored cross-language experiments by employing mBERT and XLM-RoBERTa (Conneau et al., 2020), the multilingual variants of BERT and RoBERTa, respectively.

Language Tasks: We evaluated our method on 6 linguistic tasks: suffix prediction, parts-of-speech tagging using the Penn TreeBank (Marcus et al., 1993), syntactic chunking using CoNLL 2000 shared task dataset (Tjong Kim Sang and Buchholz, 2000), CCG super-tagging using CCGBank (Hockenmaier, 2006), syntactic dependency labeling with Universal Dependencies data-set and semantic tagging using the Parallel Meaning Bank data (Abzianidze

Table 1: Example sentence with different word-level annotations.

Words	Musharraf	meets	Vajpayee	in	the	city	of	Islamabad
Suffix	-	s	-	-	-	-	-	-
POS	NP	VBZ	NP	IN	DT	NN	IN	NP
SEM	PER	ENS	PER	REL	DEF	REL	REL	GEO
Chunk	B-NP	B-VP	B-NP	B-PP	B-NP	I-NP	B-PP	B-NP
CCG	NP	((S[dcl]\NP)/PP)/NP	NP	PP/NP	NP/N	N	(NP\NP)/NP	NP

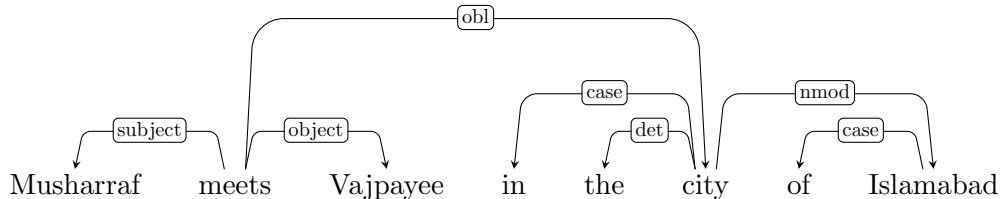


Figure 2: Syntactic relations according to the Universal Dependencies formalism. Here “Musharraf” and “Vajpayee” are the subject and object of “meets”, respectively, `obl` refers to an oblique relation of the locative modifier, `nmod` denotes the genitive relation, the prepositions “in” and “of” are treated as case-marking elements, and “the” is a determiner. See <https://universaldependencies.org/guidelines.html> for detailed definitions.

et al., 2017). See Table 1 and Figure 2 for sample annotations for different tasks. We used standard splits for training, development and test data (See Table 8 in Appendix). For multilingual experiments, we annotated a small portion of multi-parallel news data (Bojar et al., 2014) for English, German and French using RDRPOSTagger (Nguyen et al., 2014) (See Table 9 in Appendix for statistics).

Transfer Learning: For the transfer-learning experiments, the models (BERT, XLNet, and RoBERTa) were fine-tuned towards GLUE tasks. Specifically, we experimented with SST-2 for sentiment analysis with the Stanford sentiment tree-bank (Socher et al., 2013), MNLI for natural language inference (Williams et al., 2018), QNLI for Question NLI (Rajpurkar et al., 2016), RTE for recognizing textual entailment (Bentivogli et al., 2009), MRPC for Microsoft Research paraphrase corpus (Dolan and Brockett, 2005), and STS-B for the semantic textual similarity benchmark (Cer et al., 2017). All models were fine-tuned using identical settings, and we conducted three independent runs to identify the best-performing model.

Classifier Settings: We utilized a linear probing classifier with elastic-net regularization, employing a categorical cross-entropy loss function and optimizing it with Adam (Kingma and Ba, 2014). The training process involved shuffled mini-batches of size 512 and was stopped after 10 epochs. The regularization weights were trained using a grid-search algorithm. In the case of sub-word based models, we considered the last activation value as the representative of the word, as also done for the embeddings (Liu et al., 2019a; Durrani et al., 2019). Linear classifiers are commonly chosen for analyzing deep NLP models due to their enhanced interpretability (Qian et al., 2016a; Belinkov et al., 2020). Hewitt and Liang (2019) have also demonstrated that linear probes exhibit higher “Selectivity” (which

Table 2: Ablation Study: Selecting all, top, random and bottom 20% neurons and zeroing-out remaining to evaluate classifier accuracy on blind test (averaged over 3 runs). In the last column we select 20% neurons based on weight magnitude of the neurons.

	POS			SEM			Chunking		
	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa
All	96.16	96.19	96.72	92.11	92.68	93.13	95.01	94.15	95.09
Top	90.16	92.28	91.93	84.32	90.70	90.16	89.01	89.16	89.63
Random	28.45	58.17	44.56	64.28	72.14	66.15	75.83	75.26	75.30
Bottom	16.86	44.64	23.09	59.02	25.37	34.35	66.82	46.66	54.11
Wt. Mag.	67.89	64.73	71.06	61.56	68.86	64.79	68.31	75.44	73.22

we also observed), a desirable characteristic for more interpretable probes. Linear probes are especially important for our method since we employ the learned weights as a measure of each neuron’s importance.

4. Evaluation

Before presenting our analysis, we evaluate the efficacy of our neuron ranking. We validate LCA using i) ablation study, ii) classifier retraining, iii) selectivity, and iv) via qualitative evaluation.

4.1 Ablation Study

Given a classifier trained to predict a linguistic task T , we extract a ranked list of neurons and conduct ablation experiments to validate the rankings. In these experiments, we “zero-out” all activations in the test, except for the selected $M\%$ neurons. We evaluate the rankings by retaining the top, random, and bottom 20% neurons. Additionally, we establish another baseline in which we selected neurons based on their weight magnitude and ablated the remaining neurons. It is worth noting that the choice of 20% is arbitrary and not experimentally explored; it was used solely to demonstrate the effectiveness of the rankings and to select the best lambdas (see Section 2.2).

Table 2 illustrates the efficacy of our rankings. Retaining the top 20% neurons yields high prediction accuracy, in contrast to retaining the bottom or random 20% neurons. Notably, the accuracy of the **Random** neurons, although significantly lower compared to the **Top** neurons, is still relatively high in some cases (e.g., for the “Chunking” task, which is related to predicting syntax). We observed that when the underlying task is complex, the related information is more distributed across the network, leading to redundancy. The issue of redundancy in pre-trained models will be revisited in Section 5.3. This **Weight Magnitude** baseline demonstrated similar performance to randomly selecting neurons in most cases.

4.2 Minimal Neuron Subset

In the previous section, we conducted an experiment where we masked out most of the classifier neurons, leaving only 20% of them, in order to evaluate the accuracy of the rankings.

Table 3: Selecting minimal number of neurons for each downstream NLP task. Accuracy numbers reported on blind test-set (averaged over three runs) – Neu_t = Percentage of top selected neurons, Acc_a = Accuracy using all neurons, Acc_t = Accuracy using selected neurons after retraining the classifier using selected neurons, Sel = Difference between linguistic task and control task accuracy when classifier is trained on all neurons (Sel_a) and top neurons (Sel_t).

	POS			SEM			Chunking		
	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa
Neu_t	5%	5%	5%	5%	5%	5%	10%	10%	10%
Acc_a	96.16	96.19	96.72	92.11	92.68	93.13	95.09	94.15	94.68
Acc_t	95.92	96.49	96.48	92.18	92.72	92.96	94.99	94.62	94.80
Sel_a	14.49	23.53	17.93	14.08	12.98	13.33	16.30	22.77	18.39
Sel_t	31.84	31.98	31.91	27.17	26.83	27.00	29.19	28.42	28.39
Probeless	95.45	96.18	94.90	90.87	91.35	89.99	90.36	91.80	91.16

Now we leverage these rankings to select the minimum number of neurons required for each linguistic task. To achieve this, we employ Algorithm 4 to obtain a ranking of neurons. We then iteratively select a percentage, denoted as $N\%$, of the highest-ranked neurons that achieve comparable accuracy (using a threshold $\delta = 1.0$) to using the entire network (all the features). By analyzing these significant neurons, we can highlight two aspects: i) the parts of the learned network that predominantly capture different linguistic phenomena, and ii) the extent to which information is localized or distributed across different properties. A summary of the results is provided in Table 3.

Firstly, we demonstrate that for all tasks, selecting a subset of the top $N\%$ neurons and retraining the classifier can achieve similar or even better accuracy compared to using all the neurons (Acc_a) as static features. In the case of lexical tasks such as **POS** or **SEM** tagging, a very small number of neurons (approximately 5%, i.e., 499 out of 9984³) were found to be sufficient to achieve a comparable accuracy (Acc_t) to the oracle accuracy (Acc_a). On the other hand, syntactic tasks like **Chunking** tagging, which model non-local dependencies, required a larger set of neurons (999 out of 9984) to achieve the same level of accuracy. We also examined more complex syntactic tasks, namely CCG super tagging and Syntactic dependency relations, which are discussed later in the paper. These tasks required an even larger set of neurons (15%) to achieve performance close to the oracle accuracy. Additionally, we included another baseline for comparison. We employed the Probeless method (Antverg and Belinkov, 2022) to generate rankings of neurons for each task and subsequently retrained the classifier by selecting the top $N\%$ neurons based on these rankings. Our findings demonstrate that the Probeless method exhibited similar performance to our approach in the POS tagging task, while our method (LCA) outperformed Probeless in the other two tasks.

3. When training the classifier, we create a large feature vector (768×13 – hidden dimensions \times number of layers) by concatenating the layers.

4.3 Selectivity

Although selectivity is not directly used to evaluate the efficacy of our ranking, it is generally considered important to demonstrate that the probe truly reflects the knowledge contained within the understudied representation and not its ability to memorize the task. We use *Selectivity* to further demonstrate that our probes (trained using the entire representation and selected neurons) do not memorize word types but highlight the knowledge captured within the underlying representation. Recall that an effective probe is recommended to achieve high linguistic task accuracy and low control task accuracy. The results (see Table 3) show that selectivity with top neurons (Sel_t) is much higher than selectivity with all neurons (Sel_a). It is evident that using all the neurons may contribute to memorization, whereas higher selectivity with selected neurons indicates less memorization and the efficacy of our neuron selection.

4.4 Qualitative Evaluation

Visualizations have been used effectively to gain qualitative insights on analyzing neural networks (Karpathy et al., 2015; Kádár et al., 2017). We carry a qualitative evaluation and analysis of the selected neurons using the NeuroX visualization tool (Dalvi et al., 2023). We do this by visualizing: i) the activations of the top neurons w.r.t a property, ii) or analyzing the top n words that activates a neuron.

Mr. Gotlieb , who **serves** as a consultant to Stikeman ,
Elliott , one of Canada 's **biggest** law firms

From a helicopter a thousand feet above Oakland after the
deadliest earthquake in U.S. history

Traders said property-casualty companies with the **heaviest**
exposure in the San **Francisco** area include the OTC 's
Safeco and Ohio Casualty .

(a) Superlative Adjective Neuron (ROBERTA) – Layer 1: 193

It has been **trying** to improve its share of the residential
market .

He **declined** to comment on whether the company is
considering a dividend or is **planning** any acquisition .

The university is **considering** installing a \$ **250,000** system
to store applications electronically .

(b) Gerund Verb Neuron (ROBERTA) – Layer 1: 750

Figure 3: Visualizations (POS) – Neuron 193 in Layer activates positively for superlative adjectives, Neuron 750 activates positively for gerund verbs.

Figures 3-5 visualize the activations of the top neurons for a few properties. It shows how single neurons can focus on very specific linguistic properties. Figure 3a shows a neuron that activates positively for superlative adjectives (e.g. biggest, deadliest, heaviest). Figure 3b shows a neuron that activates positively for gerund verbs (e.g. trying, considering etc). Similarly Figure 4a shows a neuron that activates negatively for person name (e.g. Robert, Annan and McCain etc). Figure 4b shows a neuron that activates negatively for places, specifically countries and cities (e.g. Paris, Qatar, Doha etc). We are able to identify

Defense Secretary **Robert** Gates called him ' one of the best strategic thinkers ' in the U.S. military

Mr. **Annan** said in his monthly report to the U.N. Security Council Monday that there has been a **dangerous** increase in violence in Darfur .

McCain may announce his running mate during a rally **on** August 29 , three days before the start of the Republican National Convention .

(a) Person Name Neuron (XLNet) – Layer 2: 651

Paris is often referred to as the 'City of Light'

In the 21st century, **Qatar** emerged as a significant power in the Arab world through its resource-wealth

Over 80% of the nation's population lives in **Doha**

(b) Place Neuron (XLNet) – Layer 2: 115

Figure 4: Visualizations (SEM) – Neuron 651 in layer 2 activates negatively for person names, Neuron 115 is a place neuron.

They also violate the relevant Security Council resolutions , in particular resolution 2216 (2015) , and are consistent with the Houthis ' total rejection of the said resolution .

Figure 5: Position Neuron: Activates positively in the beginning, becomes neutral in the middle and negatively towards the end of sentence.

neurons learning a specific property. Although our results are focused on linguistic tasks, the LCA method can be used to probe any property for which labeled data can be created. For instance, Mousi et al. (2023) used LLMs to annotate encoded concepts within pre-trained language models and carried out neuron analysis using these annotations. Here, we trained a classifier to predict position of the word, i.e., identify if a given word is at the beginning, middle, or end of the sentence. As shown in Figure 5, the top neuron identified by this classifier activates with high positive value at the beginning (blue), moves to zero in the middle (white), and gets a high negative value at the end of the sentence (red).

Another way to qualitatively validate the salient neurons identified by our method is to visualize the top words that activate a given neuron. Table 4 shows salient neurons w.r.t a property in different models along with their top-5 activating words in the test-set. Looking at the top-5 activating words for a neuron, we can validate that the identified neuron indeed learns that property. These results also show how visualization can be combined with any neuron ranking method to give us a quick insight into the property a neuron has learned.

5. Analysis

Now that we have established the efficacy of our method, both empirically and qualitatively, let us analyze the discovered salient neurons. Specifically we want to study the following questions: i) how do the salient neurons distribute across the layers? ii) how do they

Table 4: Ranked list of words for some individual neurons, VBD: Past-tense verb, VBG: Gerund Verb, VBZ: Third person singular, CD: Numbers, LOC: Location, ORG: Organization, PER: Person, YOC: Year of the century.

Neuron	Property	Model	Top-5 words
Layer 9: 624	VBD	RoBERTa	supplied, deposited, supervised, paled, summoned
Layer 2: 750	VBG	RoBERTa	exciting, turning, seeing, owning, bonuses
Layer 0: 249	VBG	BERT	requiring, eliminating, creates, citing, happening
Layer 1: 585	VBZ	XLNet	achieves, drops, installments, steps, lapses, refunds
Layer 2: 254	CD	RoBERTa	23, 28, 7.567, 56, 43
Layer 5: 618	CD	BERT	360, 370, 712, 14.24, 550
Layer 1: 557	LOC	XLNet	Minneapolis, Polonnaruwa, Mwangura, Anuradhapura, Kobe
Layer 5: 343	ORG	RoBERTa	DIA, Horobets, Al-Anbar, IBRD, GSPC
Layer 10: 61	PER	RoBERTa	Grassley, Cornwall, Dalai, Bernanke, Mr.Yushchenko
Layer 6: 132	PER	BERT	Nick, Manie, Troy, Sam, Leith
Layer 2: 343	YOC	BERT	1897, 1918, 1901, 1920, Alam

distribute across the properties, iii) how redundant are the neurons w.r.t a concept? iv) how transfer learning impacts linguistic knowledge in deep NLP models? and v) how much overlap in neurons do we see in the multilingual models?

5.1 Distribution of Neurons Across Layers

Work done in interpreting deep NLP models often study the distribution of linguistic knowledge across layers by training classifiers for each layer individually (Liu et al., 2019a; Kim et al., 2020; Belinkov et al., 2020). In this study, we investigate how neurons selected from the entire network distribute across various layers of the model. This analysis provides an alternative perspective on which layers predominantly contribute to different tasks.

We explore two additional tasks in this section. A simple task of predicting the suffix of a word given its representation, and the second is a more complex task predicting CCG tag of a word. CCG captures syntactic information globally within a sentence at the word level by assigning a label to each word, indicating its syntactic role. These annotations can be conceptualized as a function that processes and produces syntactic categories (refer to Table 1 for an example).

The findings presented in Figure 6 reveal insights across various tasks and transformer models. Neurons responsible for predicting suffixes, a shallow linguistic task, are primarily concentrated in the first two layers of the models. Similarly, abstract concepts like learning morphology (POS tagging) are predominantly captured within these initial layers as well. These patterns indicate that the initial layers of the network are focused on reconstructing word structure, which is disassembled during the subword segmentation in the pre-processing stage. As the complexity of the task increases, such as learning lexical semantics, we observe that neurons in middle and higher middle layers become more active. At the opposite end of the spectrum, the intricate task of modeling syntax (CCG supertagging) is primarily handled by neurons in the final layers.

How do the architectures compare in terms of the distribution of concept neurons? We notice that while BERT adheres to the observed patterns, the top neurons in BERT are spread out more evenly across the layers compared to XLNet, where

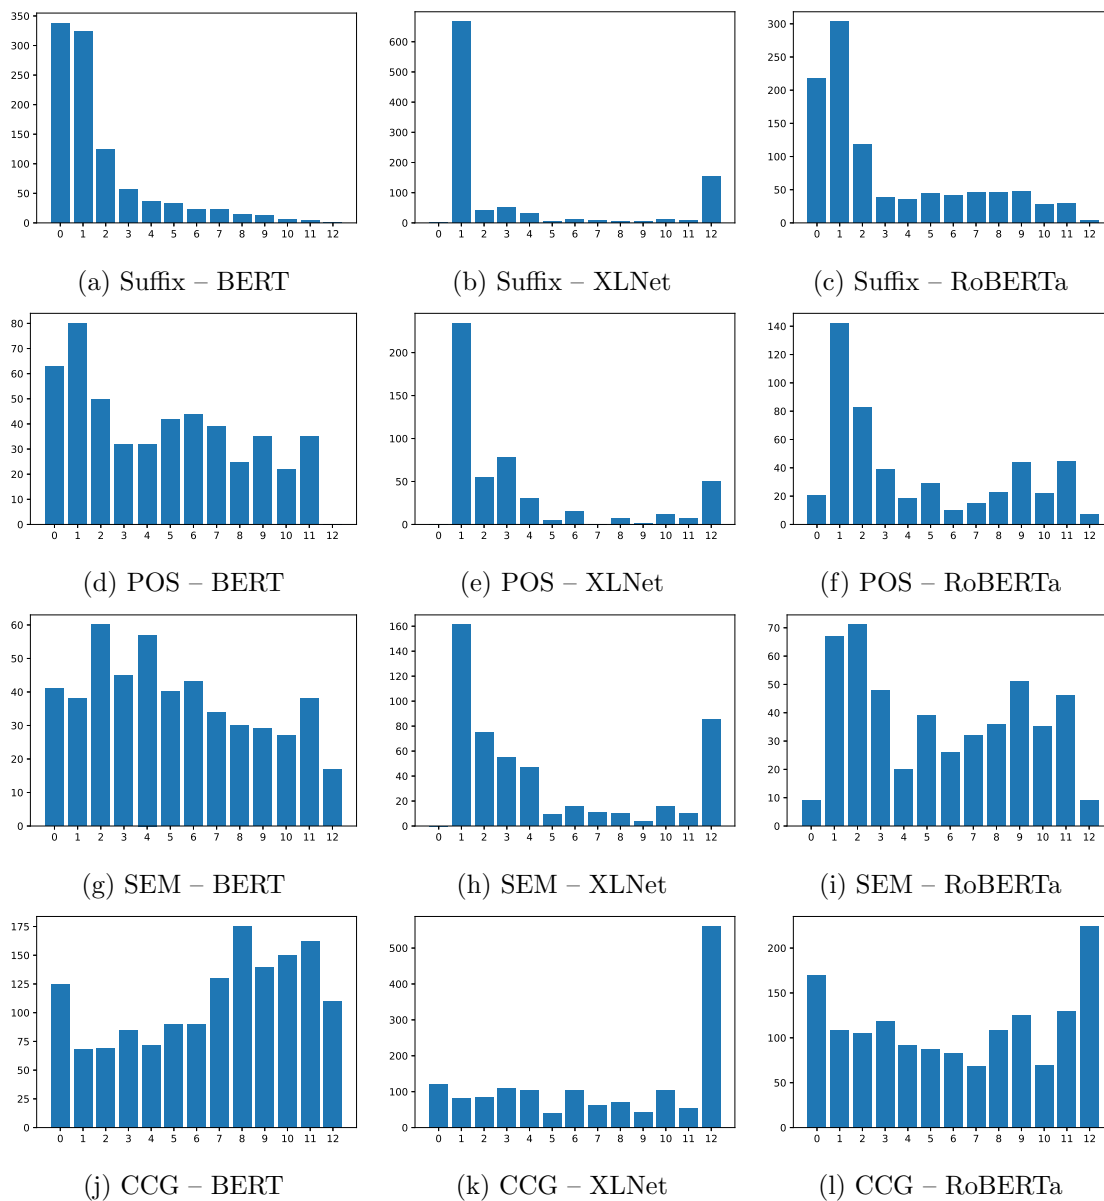


Figure 6: How top neurons distribute across different layers for each task? X-axis = Layer number, Y-axis = Number of neurons selected from that layer.

the top neurons for a specific task are concentrated in fewer layers. This indicates that in BERT, each layer contains neurons specialized in learning specific language properties, whereas in XLNet, certain designated layers are specialized for these properties. Unlike auto-encoder-based BERT models, neurons in the embedding layer consistently show minimal contribution in XLNet across all tasks. Let’s examine these distinctions for each concept:

In the case of predicting suffixes, we clearly see that the neurons capturing the information are learned in the first two layers. It is interesting to note that RoBERTa maintains this information consistently in the network. For the task of learning parts-of-speech, most layers in BERT contributed towards the top neurons, while the distribution is dominated by lower layers in XLNet, with an exception of XLNet not choosing any neurons from the embedding layer. RoBERTa exhibits similar behavior to XLNet with embedding layer contributing small number of neurons. However, in this case higher layers also contribute to the most salient neurons. On the task of semantic tagging, similar to POS, all layers of BERT contributed to the list of top neurons. However, the lower-middle layers showed the most contribution. This is in line with Liu et al. (2019a) who found middle and higher middle layers to give optimal results for the semantic tagging task. On XLNet and RoBERTa, the lower layers after the embedding layer got the largest share of the top neurons of SEM. This trend is consistent across other tasks, i.e., the core linguistic information is learned earlier in the network with an exception of BERT, which distributes information across the network.

In contrast to the lexical tasks discussed earlier, which are mainly captured in the initial and middle layers, there is a noticeable shift in pattern when it comes to learning syntactic tasks. Salient neurons for syntactic tasks emerge in the higher layers. In XLNet and RoBERTa, the majority of salient neurons were contributed by the final layer. Even in the case of BERT, which exhibits a more distributed behavior, higher layers contribute more neurons than the initial layers. These findings align with previous research Liu et al. (2019a); Tenney et al. (2019) on analyzing contextualized representations across layers to explore how various linguistic concepts are learned at different layers. Our neuron analysis method serves as a complementary approach to their study, reinforcing the overall understanding of the layer-wise distribution of linguistic knowledge in transformer models.

5.2 Distribution of Neurons Across Properties

The Linguistic Correlation Analysis provides a ranking for a property set (e.g., Chunking tagging) as well as for individual labels (e.g., Chunking:B-VP). The ranking algorithm extracts neurons for each label l in task T , sorted based on absolute weights. The final rankings are obtained by selecting neurons from each label using the neuron ranking algorithm as described in Section 2.3. The distribution of neurons, allows us to analyze how localized or distributed a property is, based on the number of neurons that are selected for each label in the task. Figure 7 shows various properties alongside the number of neurons (in the minimal set) required to capture each property. We observed that functional words, also known as closed class categories such as determiners or numbers, are localized in as few as 10 neurons. In contrast, open class words like nouns and verbs are distributed across a larger number of neurons. Open class categories occur in a variety of contexts, necessitating more neurons to capture diverse linguistic nuances. This observation held consistently across different properties within various tasks, as well as across different models and languages. For detailed results, please refer to Appendix 7.

How do individual properties distribute across layers? Previously, our analysis focused on examining each linguistic task as a whole. Now, we study the individual properties (e.g., adjectives) to understand whether they are localized or widely distributed across lay-

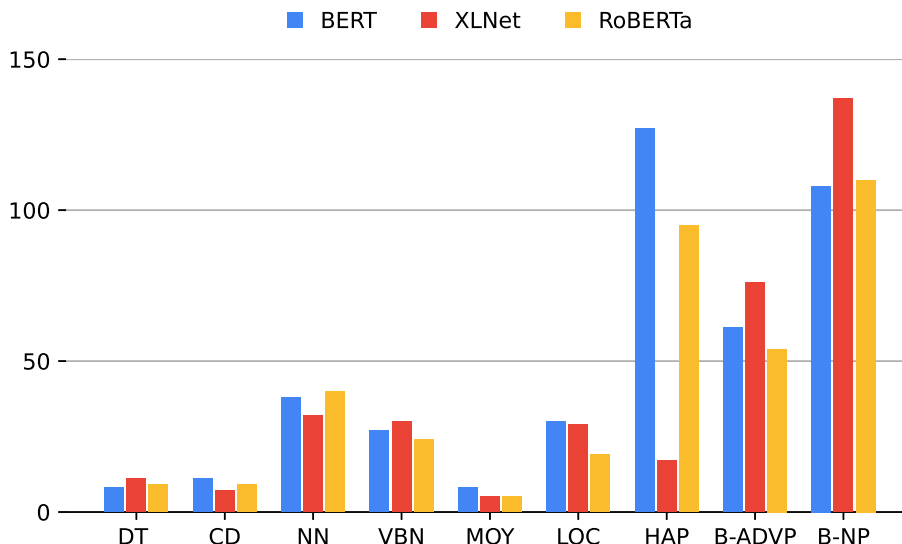


Figure 7: Localized vs. Distributed Properties: DT = Determiner, CD = Numbers, NN = Noun, VBN = Past participle verb, MOY, Month of Year, LOC = Location, HAP = Event (happening), B-ADVP = Beginning of Adverb Phrase, B-NP = Beginning of Noun Phrase, Y-axis = number of neurons assigned to this property

ers in different architectures. We observed intriguing similarities across architectures. For instance, neurons predicting the foreign words (**FW**) property were predominantly localized in the final layers (BERT: 13, XLNet: 11, RoBERTa: 11) in the architectures under investigation. In contrast, neurons capturing common class words such as adjectives (**JJ**) and locations (**LOC**) were localized in lower layers (BERT: 0, XLNet: 1, RoBERTa: 1), indicating that these common class words might be determined at a lexical level. However, we did find variations in certain cases. For example, personal pronouns (**PRP**) or plural nouns (**NNS**) in POS tagging and event class (**EXC**) in semantic tagging were handled at different layers across different architectures.

5.3 Redundancy

Neural networks are inherently distributed and, as a result, exhibit natural redundancy. Pretrained models, due to over-parameterization and the incorporation of techniques like drop-outs, further contribute to this redundancy. A growing body of research suggests that these models don’t utilize the full representational capacity offered by their complex architectures during inference (Michel et al., 2019). This hypothesis has motivated researchers to prune components during inference, a process that has been shown to not significantly impact model performance (Voita et al., 2019; Sajjad et al., 2020).

In Section 4.2, we identified the most salient neurons for various linguistic tasks. Now we will explore whether the remaining neurons in the network contain redundant information. To conduct this experiment, we randomly selected N neurons from the model and retrained a classifier to predict specific properties. We set N equal to the number of top neurons selected for the task. The results presented in Table 5 demonstrate that using the same

Table 5: Selecting minimal number of neurons for each downstream NLP task. Accuracy numbers reported on blind test-set – Neu_t = Percentage of top selected neurons, Acc_a = Accuracy using all neurons, Acc_t = Accuracy using selected neurons after retraining the classifier using selected neurons, Acc_r = Accuracy using randomly selected neurons after retraining the classifier, Acc_b = Accuracy using bottom neurons after retraining the classifier.

	POS			SEM			Chunking		
	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa
Neu_t	5%	5%	5%	5%	5%	5%	10%	10%	10%
Acc_a	96.16	96.19	96.72	92.11	92.68	93.13	95.09	94.15	94.68
Acc_t	95.92	96.49	96.48	92.18	92.72	92.96	94.99	94.62	94.80
Acc_r	94.12	95.19	94.38	90.01	90.37	91.17	92.57	91.63	93.10
Acc_b	93.20	91.20	93.21	90.00	88.80	90.70	91.60	92.35	92.11

number of randomly selected neurons yields very similar performance (within a margin of less than 3%) compared to the classifier trained with the top neurons. Surprisingly, even the least significant neurons carry sufficient information to predict the tasks accurately. This aligns with the findings of [Prasanna et al. \(2020\)](#), who revealed the high re-trainability of even the poorest sub-networks in BERT. These results have practical implications for network pruning and efficient feature-based transfer learning [Dalvi et al. \(2020\)](#).

5.4 Transfer Learning

Continuing our neuron analysis of linguistic knowledge learned within transformer-based language models, we dive into further experiments employing our LCA method to investigate the redistribution of morpho-syntactic and semantic knowledge within these pretrained models as they undergo fine-tuning for GLUE tasks ([Wang et al., 2018](#)). Specifically, we focus on extracting the most salient neurons related to specific linguistic properties (e.g., POS task), from both the pre-trained model and its fine-tuned counterpart. By comparing the distribution of these neurons across the network, we gain valuable insights into how linguistic knowledge evolves during the fine-tuning process. Below we discuss notable findings from this set of experiments:

Information becomes less distributed in the fine-tuned XLNet and RoBERTa models post fine-tuning. Table 6 presents the classifier accuracy results obtained by selecting the most (top) and least (bottom) N% salient neurons across various tasks and models. Comparing the performance of these neurons in fine-tuned models to the baseline model, a significant decrease in accuracy is evident, particularly in the cases of RoBERTa and XLNet. This decline implies that the information stored in the baseline models is more redundant, given that the bottom neurons still retain linguistic knowledge, as indicated in Section 5.3 on redundancy. In contrast, the fine-tuned models exhibit more localized and less distributed information, a consequence of their specific adaptation for downstream NLP tasks. Noteworthy is the finding that the bottom neurons in the fine-tuned BERT model

Table 6: Minimal top and bottom number of neurons for each downstream NLP task before (Baseline) and after fine-tuning. Task_t = Accuracy using top selected neurons after retraining the classifier using selected neurons, Task_b = Accuracy using bottom neurons after retraining the classifier using selected neurons.

	POS			SEM			Chunking		
	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa	BERT	XLNet	RoBERTa
Baseline _t	95.92	96.49	96.48	92.18	92.72	92.96	94.99	94.62	94.80
Baseline _b	93.20	91.20	93.21	90.00	88.80	90.70	91.60	92.35	92.11
MRPC _t	95.91	95.24	95.23	92.2	92.33	91.53	94.97	93.8	94.5
MRPC _b	93.13	78.83	91.11	90.04	72.33	88.14	91.92	88.0	91.9
QNLI _t	96.09	94.71	95.81	92.09	91.41	91.54	94.33	93.99	94.08
QNLI _b	93.47	10.39	84.31	90.01	17.84	78.67	91.78	79.37	90.71
MNLI _t	95.81	94.93	95.45	91.9	91.34	91.43	93.81	92.86	93.24
MNLI _b	93.03	41.13	84.85	89.97	43.05	79.02	91.33	85.02	89.53
RTE _t	95.93	95.24	95.63	92.0	91.36	91.56	94.89	94.01	94.33
RTE _b	93.11	87.96	90.42	90.03	80.01	86.94	92.01	89.98	90.91
SST _t	95.92	95.02	95.67	92.1	91.33	91.51	94.48	93.74	94.17
SST _b	93.38	35.22	60.43	89.93	34.83	60.73	91.91	80.34	82.86
STS _t	95.94	95.15	95.76	92.1	91.68	91.52	94.89	93.82	94.31
STS _b	93.12	88.14	85.91	90.03	83.78	79.56	92.01	89.93	88.89

experienced minimal change, suggesting that linguistic information remains redundant and distributed in BERT even after the fine-tuning process.

How do salient neurons distribute across the network layers? In the earlier section (Section 5.1), we explored the distribution of salient neurons across network layers for each task. Now, our focus shifts to examining how this distribution changes in the fine-tuned models. The results for selected GLUE tasks are depicted in Figure 8.⁴ The graph illustrates the number of neurons across layers in both the baseline model and each fine-tuned model for different downstream NLP tasks. Notably, in RoBERTa and XLNet, the most salient neurons are now found in the lower layers. This is particularly striking in *Roberta-SST* and *XLNet-QNLI* (refer to Figures 8f and 8e), where the number of salient chunking neurons significantly increases in the lower layers while decreasing in the higher layers compared to the baseline. These observations signify a trend where neurons in lower layers bear more responsibility, focusing on learning task-specific information, while the higher layers become specialized in this regard.⁵ In contrast, BERT does not exhibit this behavior. Unlike XLNet, where linguistic properties localize in the lower layers, BERT distributes information evenly across the network, as previously noted in Section 5.1, where we found linguistic properties in XLNet to be localized in the lower layers, while BERT disperses information throughout all layers. Our neuron analysis proves valuable in understanding the essential linguistic

4. For a comprehensive list of tasks and linguistic properties, please refer to the Figure 12 in the Appendix.

5. Rangamani et al. (2023) recently demonstrated how neural collapse pushes structured features to the lower layers of the network. Durrani et al. (2022) made similar observations when studying how latent spaces evolve as the models are fine-tuned towards downstream NLP tasks.

concepts for downstream tasks. A more detailed exploration would involve pinpointing the specific concepts regressed to lower layers versus those retained in higher layers. However, this in-depth analysis remains a topic for future research.

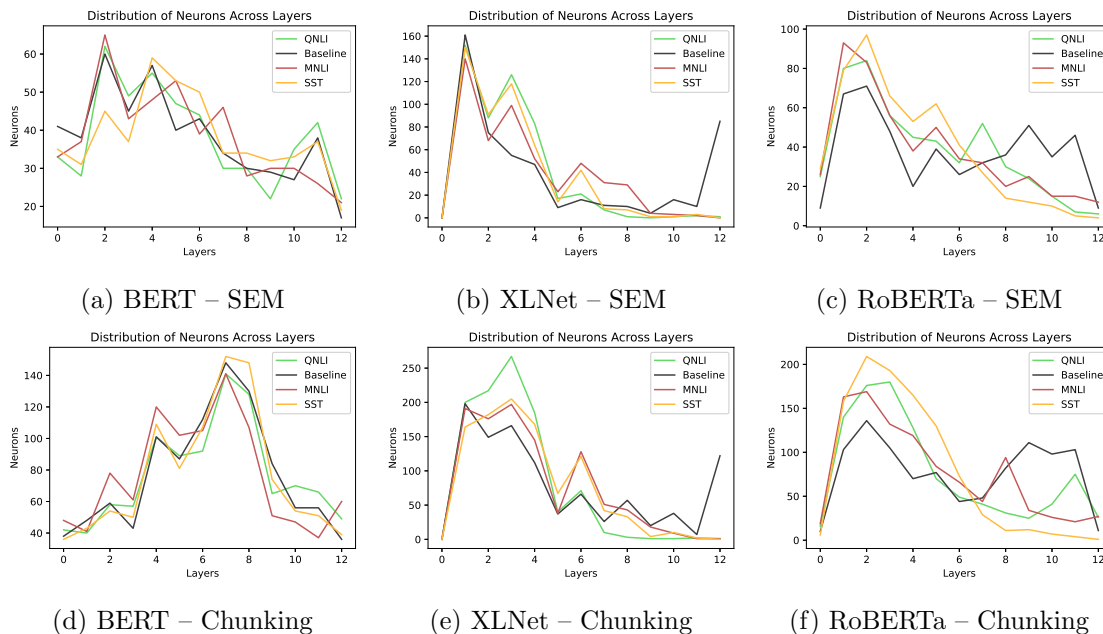


Figure 8: Distribution of top neurons across layers. Salient linguistic neurons regress towards lower layers in XLNet and RoBERTa models as the models are fine-tuned towards GLUE tasks.

5.5 Multilingual Models

Recently, there has been growing interest in cross-lingual transfer learning, particularly in scenarios transitioning from high-resource to zero-shot or low-resource contexts. In this study, we investigate the effectiveness of multilingual BERT (Devlin et al., 2019) and XLM-RoBERTa (Conneau et al., 2020) using our Linguistic Correlation Analysis. We revisit our analysis focusing on parts-of-speech tagging, a task for which annotations are readily available in multiple languages. Additionally we use syntactic dependency labeling, the task of assigning label to an arc that connects two words in a syntactic dependency tree. A dependency is a directed bi-lexical relation between a head and its dependent, or modifier. We simply concatenate the activations of the words in question and train classifiers on the concatenated vectors.⁶ We experiment with English, German and French. We used monolingual models including FlauBERT (Le et al., 2020) and BERT-base-German-cased (Chan et al., 2020) for French and German respectively. Below we discuss notable findings from this set of experiments:

6. We also explored variations i) using just the activation of the head or modifier word, ii) use average vector of the two. Both variations resulted in sub-optimal classifier accuracy compared to using concatenation.

Table 7: Selecting salient neurons for POS and Syntactic Dependency Labeling tasks. Accuracy numbers reported on blind test-set (averaged over three runs) – Acc_a = Accuracy using all neurons, $Acc_{t,r,b}$ = Accuracy using top, bottom, and randomly selected neurons after retraining the classifier using selected neurons.

	POS English			POS German			POS French		
	BERT	mBERT	XLM-R	BERT	mBERT	XLM-R	BERT	mBERT	XLM-R
Acc_a	94.82	94.48	94.70	93.37	93.50	93.91	96.28	95.84	96.22
Acc_t	94.52	94.60	94.17	93.82	93.42	93.77	96.24	95.59	95.39
Acc_r	93.96	94.55	94.00	93.40	93.30	93.63	95.50	95.31	95.43
Acc_b	93.52	94.48	92.20	93.32	93.01	92.04	95.04	95.04	94.27
	Dependency English			Dependency German			Dependency French		
	BERT	mBERT	XLM-R	BERT	mBERT	XLM-R	BERT	mBERT	XLM-R
Acc_a	91.90	91.50	92.14	91.56	91.45	92.00	93.52	94.30	94.62
Acc_t	91.18	91.52	92.05	91.62	91.35	92.06	92.86	94.33	94.53
Acc_r	90.90	91.10	91.46	90.77	90.87	91.73	92.62	93.81	94.06
Acc_b	88.01	88.00	89.94	87.67	87.18	90.18	91.09	91.81	92.55

Salient Neurons: We found that a subset of neurons (5% for POS tagging as before) can be extracted to obtain comparable accuracy (Acc_t) to the oracle (Acc_a) which uses the entire network (See Table 7). For a more complex task of predicting dependency relation given head and modifier word, a subset of 15% neurons was found to be sufficient.

Redundancy: Despite the fact that multilingual models are designed to learn multiple languages simultaneously, and the encoding space is shared, these models exhibited significant redundancy, similar to their monolingual counterparts. It is notable how random and bottom N% neurons retain retrainable information crucial for restoring the majority of the performance. Although this redundancy enhances the model’s robustness, it also complicates the task of isolating and controlling specific concepts within the network.

Layer-wise Analysis: Next, we analyze the distribution of salient neurons across layers. Our analysis reveals several findings: i) multilingual models also capture morphology in the lower layers of the model (See Figure 9), ii) the neurons that capture syntactic dependency between a head and its modifier are predominantly found at layers 4 and 5 across different languages and models (See Figure 10), iii) the neuron distribution pattern for multilingual models is quite similar across the three languages. However, monolingual models exhibit slightly different neuron distributions across the three languages. This difference may arise due to the models being trained with different data and optimization functions. For example, FlauBERT only uses the MLM (masked language model) objective and does not include NSP (next sentence prediction) during training. Similarly, the German BERT model is trained with the WWM (whole word masking) optimization function. Such variations can cause minor differences, but the general pattern of morphology captured in the lower-middle part of the network remains consistent across languages and different models.

Property-wise Analysis: Lastly, we analyze the distribution of neurons across individual properties within the task. To carry out this analysis, we converted POS tags for different

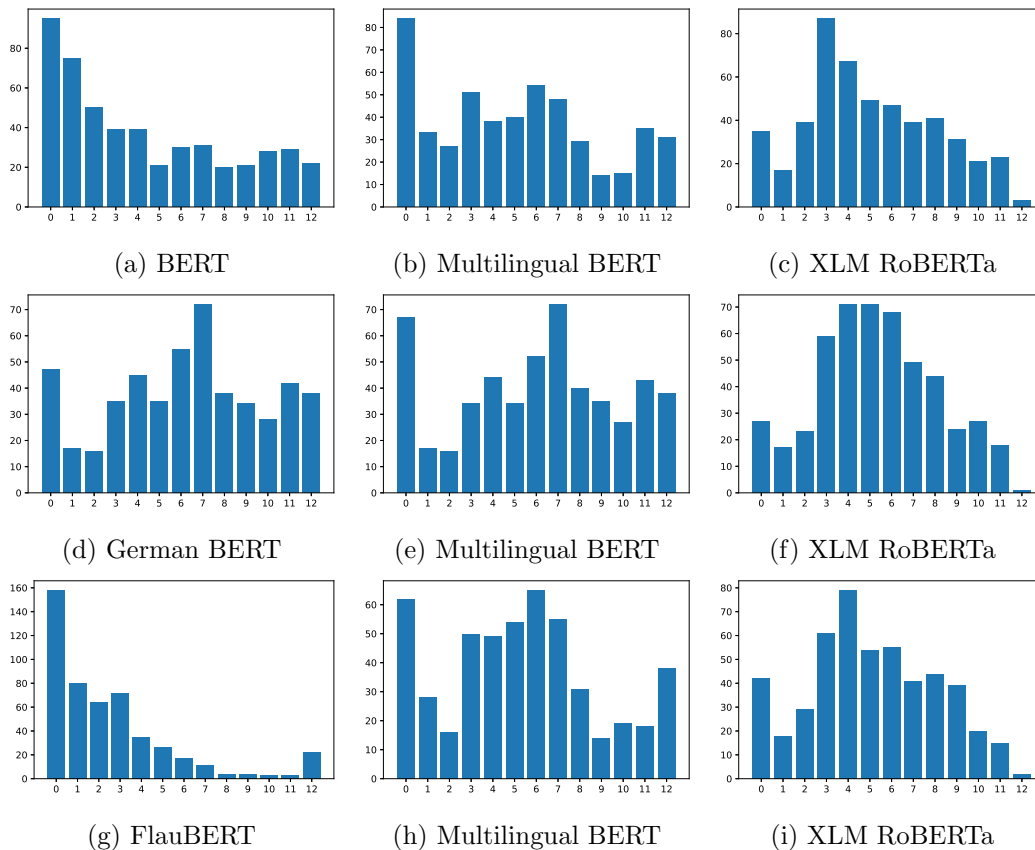


Figure 9: How salient neurons in POS task distribute across different layers for different models? English (a,b,c), German (d,e,f), French (g,h,i), X-axis = Layer number, Y-axis = Number of neurons selected from that layer.

languages into universal tags using the mapping described in [Petrov et al. \(2012\)](#). The syntactic dependency labels are shared because we used Universal dependencies. Figure 11 shows the results for the multilingual BERT model.⁷ We observe that different properties in the POS tagging task exhibit very consistent neuron distribution patterns across three different languages. Additionally, we found that functional properties such as determiners (DET) and conjunctions (CONJ) are localized to fewer neurons, while open categories such as nouns and verbs require a larger number of neurons to be accurately represented in the multilingual models. In comparison to POS tagging, where we observed harmony among languages, we found that a significantly larger number of neurons were necessary for German to capture different properties in the task of predicting syntactic dependency relations. This can be attributed to the fact that German has a more complex word order. For example German has a complicated verb complex often called as “separable verbs” that are split across the sentence, requiring the model to capture long distance dependency ([Fraser et al., 2013](#)). Similarly note how neurons that capture compound category is much

7. Please see the Appendix for results on XLM-Roberta (see Figure 14).

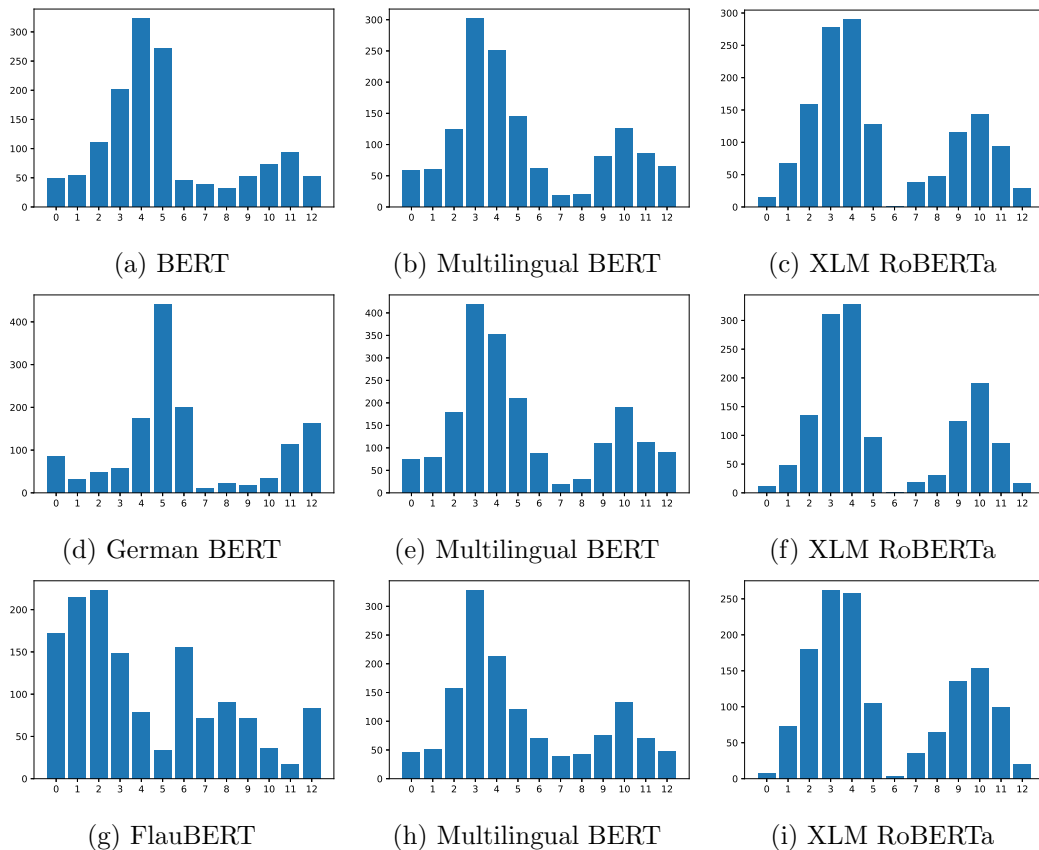
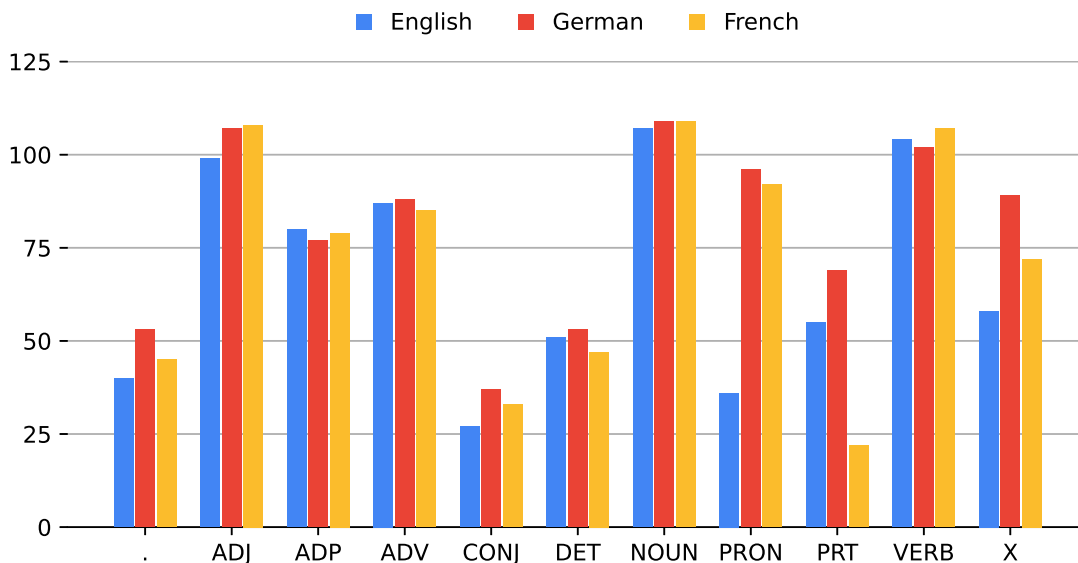


Figure 10: How salient neurons in Syntactic Dependency task distribute across different layers for different models? English (a,b,c), German (d,e,f), French (g,h,i), X-axis = Layer number, Y-axis = Number of neurons selected from that layer.

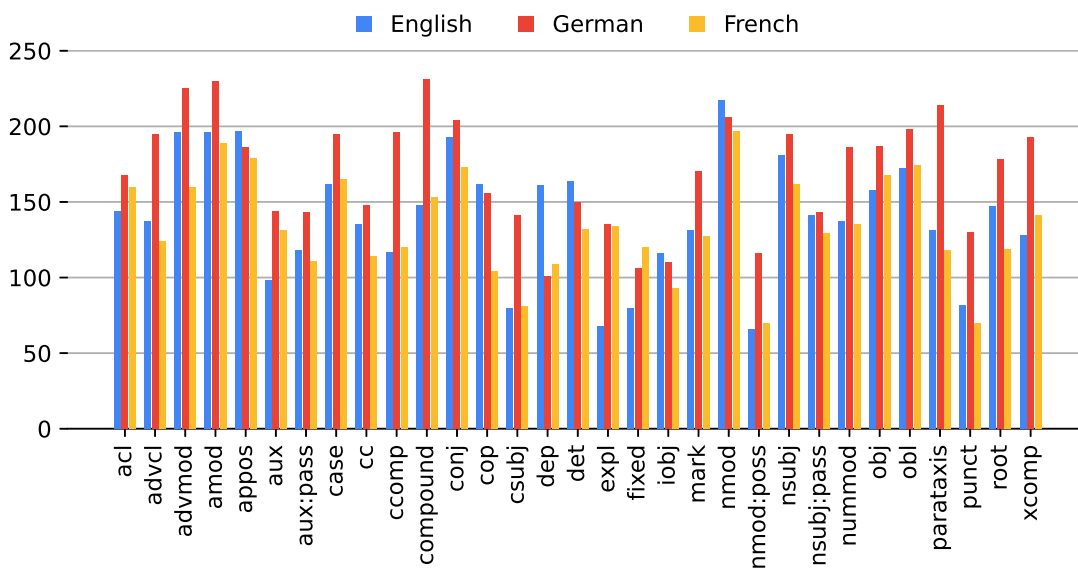
higher in German where compounding is very productively used in forming new words (Koehn and Knight, 2003).

6. Related Work

A multitude of papers have delved into the interpretation of deep NLP models, with a focus on two key questions: i) comprehending the acquired knowledge within contextualized representations, known as *Concept Analysis*, and ii) exploring how this information is utilized in the decision-making process, referred to as *Attribution Analysis*. The former entails post-hoc decomposability, where researchers analyze representations to unveil linguistic and non-linguistic phenomena captured during the network’s training for any NLP task (Belinkov et al., 2017a; Conneau et al., 2018; Sajjad et al., 2022). On the other hand, the latter focuses on characterizing the contributions of model components and input features toward specific predictions (Linzen et al., 2016b; Gulordava et al., 2018; Marvin and Linzen, 2018). The work done on neuron analysis predominantly falls under *Concept Analysis*. We classify these into a group of three:



(a) POS tagging



(b) Syntactic Dependency Labeling

Figure 11: Distribution of Top Neurons Across POS (a) and Syntactic Dependency Properties in Multilingual BERT

Visualizing provides a direct approach to understanding neuron function by identifying patterns across a set of sentences manually (Li et al., 2016; Karpathy et al., 2015). Yet, due to the sheer number of neurons in a neural network model, visualizing all of them becomes unwieldy. To address this challenge, specific cues have been employed to narrow down the selection of neurons for visualization. These cues include identifying saturated neurons, high/low variance neurons, or omitting dead neurons (Karpathy et al., 2015) when

utilizing the ReLU activation function. Despite the effectiveness of visualization in neuron analysis, it does have limitations. Firstly, it is a qualitative technique. Secondly, visualizing neurons with multiple meanings (polysemous neurons) proves challenging. Lastly, not all neurons can be visually interpreted. However, by employing a neuron ranking method such as LCA, as demonstrated in this paper, we can enhance the efficiency of qualitative analysis by focusing on the most prominent neurons.

Corpus-based Methods investigate the role of a neuron using techniques such as ranking sentences in a corpus, generating synthetic sentences to maximize its activation, or computing various neuron-level statistics over a corpus. Kádár et al. (2017) utilized corpus ranking to identify words that activate a neuron maximally by analyzing the top-k context from a corpus. In contrast, Na et al. (2019) introduced lexical concepts of different granularity, such as words, phrases, and sentences, using parse trees where each node represented a concept. However, one limitation of corpus ranking is its confined analysis space, limited to the underlying corpus. To overcome this limitation, Poerner et al. (2018) generated novel sentences maximizing a neuron’s activation, providing a broader understanding of the neuron’s role in the model. Moreover, Mu and Andreas (2020) and Suau et al. (2020) proposed a *Masking-based* method to identify important neurons for a concept. The concept involves identifying the overlap between the presence of a concept in a set of sentences and the high activation values of a neuron. Additionally, Antverg and Belinkov (2022) used data activations to derive a neuron ranking for a concept class, comparing differences in activation values across different concepts within the class. Corpus-based methods address some limitations of the probing framework inherited by LCA. However, unlike LCA, they do not model the selection of groups of neurons collaborating in learning a concept.

Classifiers Closer to our work is Torroba Hennigen et al. (2020) and Stańczak et al. (2023). They train a generative classifier with the assumption that neurons exhibit a Gaussian distribution. They fit a multivariate Gaussian over all neurons and extracted individual probes for single neurons. A caveat to their approach is that activations do not always follow a Gaussian prior in practice – hence restricting their analysis to only the neurons that satisfy this criteria. Chowdhury et al. (2024) used the LCA to carry out neuron-wise analyses of pretrained speech models.

Unsupervised Methods The literature offers a diverse array of unsupervised techniques. Researchers such as Lakretz et al. (2019) and Li et al. (2016) have employed *Neuron Ablation* to pinpoint the most significant neurons in the network. However, a drawback of single neuron ablation is that the resulting ranking may lack meaningful interpretation in large neural networks. Ablating neurons in groups and exploring all permutations or a subset of neurons becomes an NP-hard problem (Binshtok et al., 2007). Taking a different approach, Valipour et al. (2019) utilized *Random Forests* to establish a neuron ranking. Alammar (2020) applied *Matrix Factorization* (Olah et al., 2018) to decompose a large matrix into smaller matrices of factors, where each factor represented a group of elements performing a similar function. Similarly, Meyes et al. (2020) utilized UMAP (McInnes et al., 2020) to project neuron activations into a lower-dimensional space and subsequently applied K-means clustering to group the neurons. Additionally, Bau et al. (2019) conducted a multi-model search to identify the most salient neurons in the network. They employed *Pearson*

correlation to compute a similarity score for each neuron in a model concerning the neurons in other models.

Attribution Neurons The work on analyzing how knowledge is encoded within the learned representations, does not address the question whether it is used by the model during prediction is a less explored area (Feder et al., 2021; Elazar et al., 2021). There is however, another body of work called *Attribution Analysis* that focuses on the methods that characterize the role of neurons and layers towards a specific prediction. Some of the notable methods help us to identify important neurons (w.r.t prediction) in different layers of a deep neural network, include gradient and perturbation-based attribution algorithms such as Integrated Gradients (Sundararajan et al., 2017), Layer Conductance (Dhamdhare et al., 2018), Saliency (Simonyan et al., 2014) and SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017). More recent and advanced attribution algorithms that take feature or neuron interactions into account include Integrated Hessians (Janizek et al., 2020), Shapely Taylor index (Dhamdhare et al., 2020) and Archipelago (Tsang et al., 2020). However, these methods do not provide linguistic explanation for the neurons. By connecting concept analysis with the attribution analysis we can bridge this gap. We leave this exploration for the future.

7. Conclusion

In this article, we presented *Linguistic Correlation Analysis* to discover salient neurons in deep NLP models. We train linear probes with elastic-net regularization to extract a neuron ranking of a pre-trained model, with respect to any extrinsic property of interest. We discovered and analyzed individual neurons across a variety of neural language models on the task of predicting core linguistic properties (morphology, syntax and semantics). Our results reinforce previous findings and also illuminate further insights:

- Although neural language models contain highly distributed information, it is feasible to distill a limited set of neurons to perform subsequent natural language processing tasks.
- The number of salient neurons depends on the complexity of the property in question. For instance, functional classes like closed-class words tend to be concentrated in a smaller number of neurons, while open-class words are spread across a larger neural network.
- Information is redundantly dispersed, with multiple subsets of neurons carrying identical information coexisting within the network.
- Neurons specializing in acquiring basic lexical details like suffixation or word structure are concentrated in the lower layers of the model, whereas those capturing non-local dependencies are primarily situated in the middle and higher layers of the network.
- During transfer learning, the neurons responsible for learning linguistic information are regressed to lower layers in the model, as the higher layers are optimized for specific downstream tasks.

- Fascinating distinctions between architectures were noticed: in XLNet, neurons were concentrated within fewer layers, whereas in BERT, they remained spread throughout the network even after fine-tuning.
- Our research revealed that multilingual models demonstrated a comparable distribution of neurons when learning concepts across various languages.

Future Work Most of the work done on neuron analysis including ours only highlights the knowledge encoded within the learned representations. The question whether it is actually used by the model during prediction has been largely unanswered. A few recent studies carried to answer this question such as Feder et al. (2021) and Elazar et al. (2021) analyze representations at a more holistic level and ignore individual neurons. More recently Antverg and Belinkov (2022) showed that high probing accuracy does not necessarily mean that the information is important by the model. In the future, we would try to establish connection between encoded knowledge versus how it is utilized during prediction to enable richer explanations.

References

- Lasha Abzianidze, Johannes Bjerva, Kilian Evang, Hessel Haagsma, Rik van Noord, Pierre Ludmann, Duc-Duy Nguyen, and Johan Bos. The parallel meaning bank: Towards a multilingual corpus of translations annotated with compositional meaning representations. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, EACL '17, pages 242–247, Valencia, Spain, 2017.
- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. Fine-grained Analysis of Sentence Embeddings Using Auxiliary Prediction Tasks. *arXiv preprint arXiv:1608.04207*, 14(14), 2016.
- J Alammar. Interfaces for explaining transformer language models, 2020. URL <https://jalamar.github.io/explaining-transformers/>.
- Omer Antverg and Yonatan Belinkov. On the pitfalls of analyzing individual neurons in language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=8uz0EWPQIMu>.
- Anthony Bau, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. Identifying and controlling important neurons in neural machine translation. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=H1z-PsR5KX>.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. What do Neural Machine Translation Models Learn about Morphology? In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, Vancouver, July 2017a. Association for Computational Linguistics. URL <https://aclanthology.coli.uni-saarland.de/pdf/P/P17/P17-1080.pdf>.

- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. Evaluating Layers of Representation in Neural Machine Translation on Part-of-Speech and Semantic Tagging Tasks. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (IJCNLP)*, November 2017b.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. On the linguistic representational power of neural machine translation models. *Computational Linguistics*, 45(1):1–57, March 2020.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. The fifth pascal recognizing textual entailment challenge. In *In Proc Text Analysis Conference (TAC'09)*, 2009.
- Maxim Binshtok, Ronen I Brafman, Solomon Eyal Shimony, Ajay Martin, and Crag Boutilier. Computing optimal subsets. In *Proceedings of the Twenty Second AAAI Conference on Artificial Intelligence (AAAI, Oral presentation)*, July 2007.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amant, Radu Soricut, Lucia Specia, and Aleš Tamchyna. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA, June 2014. Association for Computational Linguistics. doi: 10.3115/v1/W14-3302. URL <https://aclanthology.org/W14-3302>.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/S17-2001. URL <https://www.aclweb.org/anthology/S17-2001>.
- Branden Chan, Stefan Schweter, and Timo Möller. German’s next language model. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6788–6796, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-main.598. URL <https://aclanthology.org/2020.coling-main.598>.
- Shammur Absar Chowdhury, Nadir Durrani, and Ahmed Ali. What do end-to-end speech models learn about speaker, language and channel information? a layer-wise and neuron-level analysis. *Computer Speech and Language*, 83:101539, 2024. ISSN 0885-2308. doi: <https://doi.org/10.1016/j.csl.2023.101539>. URL <https://www.sciencedirect.com/science/article/pii/S088523082300058X>.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, July 2018.

- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale, 2020.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, and Stephan Vogel. Understanding and Improving Morphological Learning in the Neural Machine Translation Decoder. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (IJCNLP)*, November 2017.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, D. Anthony Bau, and James Glass. What is one grain of sand in the desert? analyzing individual neurons in deep nlp models. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI, Oral presentation)*, January 2019.
- Fahim Dalvi, Hassan Sajjad, Nadir Durrani, and Yonatan Belinkov. Analyzing redundancy in pretrained transformer models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP-2020)*, Online, November 2020. Association for Computational Linguistics.
- Fahim Dalvi, Nadir Durrani, and Hassan Sajjad. Neurox library for neuron analysis of deep nlp models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 75–83, Toronto, Canada, July 2023. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, Minneapolis, Minnesota, 2019. Association for Computational Linguistics.
- Kedar Dhamdhere, Mukund Sundararajan, and Qiqi Yan. How important is a neuron?, 2018.
- Kedar Dhamdhere, Ashish Agarwal, and Mukund Sundararajan. The shapley taylor interaction index, 2020.
- William B. Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*, 2005. URL <https://www.aclweb.org/anthology/I05-5002>.
- Nadir Durrani, Fahim Dalvi, Hassan Sajjad, Yonatan Belinkov, and Preslav Nakov. One size does not fit all: Comparing NMT representations of different granularities. 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 1504–1516, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1154. URL <https://www.aclweb.org/anthology/N19-1154>.

- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Firoj Alam. On the transformation of latent space in fine-tuned nlp models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1495–1516, Abu Dhabi, UAE, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.395. URL <https://aclanthology.org/2022.emnlp-main.97>.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. Amnesic probing: Behavioral explanation with amnesic counterfactuals. *Transactions of the Association for Computational Linguistics*, 9:160–175, 2021. doi: 10.1162/tacl.a.00359. URL <https://aclanthology.org/2021.tacl-1.10>.
- Amir Feder, Nadav Oved, Uri Shalit, and Roi Reichart. CausaLM: Causal model explanation through counterfactual language models. *Computational Linguistics*, 47(2):333–386, June 2021. doi: 10.1162/coli.a.00404. URL <https://aclanthology.org/2021.cl-2.13>.
- Alexander Fraser, Helmut Schmid, Richárd Farkas, Renjing Wang, and Hinrich Schütze. Knowledge sources for constituent parsing of German, a morphologically rich and less-configurational language. *Computational Linguistics*, 39(1):57–85, March 2013. doi: 10.1162/COLI.a.00135. URL <https://aclanthology.org/J13-1005>.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. Colorless green recurrent networks dream hierarchically. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1108. URL <https://www.aclweb.org/anthology/N18-1108>.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer New York, 2013. ISBN 9781489905185. URL <https://books.google.com.qa/books?id=DdQDmwEACAAJ>.
- John Hewitt and Percy Liang. 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, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1275. URL <https://www.aclweb.org/anthology/D19-1275>.
- Julia Hockenmaier. Creating a CCGbank and a wide-coverage CCG lexicon for German. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, ACL ’06*, pages 505–512, Sydney, Australia, 2006.
- Dieuwke Hupkes, Sara Veldhoen, and Willem H. Zuidema. Visualisation and ‘diagnostic classifiers’ reveal how recurrent and recursive neural networks process hierarchical structure. *CoRR*, abs/1711.10203, 2017. URL <http://arxiv.org/abs/1711.10203>.
- Joseph D. Janizek, Pascal Sturmfels, and Su-In Lee. Explaining explanations: Axiomatic feature interactions for deep networks, 2020.

- Ákos Kádár, Grzegorz Chrupała, and Afra Alishahi. Representation of linguistic form and function in recurrent neural networks. *Computational Linguistics*, 43(4):761–780, December 2017. doi: 10.1162/COLI_a.00300. URL <https://www.aclweb.org/anthology/J17-4003>.
- Andrej Karpathy, Justin Johnson, and Li Fei-Fei. Visualizing and understanding recurrent networks, 2015.
- Taeuk Kim, Jihun Choi, Daniel Edmiston, and Sang goo Lee. Are pre-trained language models aware of phrases? simple but strong baselines for grammar induction, 2020.
- Diederik Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 14(14), 2014.
- Philipp Koehn and Kevin Knight. Empirical methods for compound splitting. In *10th Conference of the European Chapter of the Association for Computational Linguistics*, Budapest, Hungary, April 2003. Association for Computational Linguistics. URL <https://aclanthology.org/E03-1076>.
- Yair Lakretz, German Kruszewski, Theo Desbordes, Dieuwke Hupkes, Stanislas Dehaene, and Marco Baroni. The emergence of number and syntax units in LSTM language models. 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 11–20, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1002. URL <https://www.aclweb.org/anthology/N19-1002>.
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Allauzen, Benoit Crabbé, Laurent Besacier, and Didier Schwab. FlauBERT: Unsupervised language model pre-training for French. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2479–2490, Marseille, France, May 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.302>.
- Piyawat Lertvittayakumjorn, Lucia Specia, and Francesca Toni. FIND: Human-in-the-Loop Debugging Deep Text Classifiers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 332–348, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.24. URL <https://aclanthology.org/2020.emnlp-main.24>.
- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. Visualizing and understanding neural models in NLP. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 681–691, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1082. URL <https://www.aclweb.org/anthology/N16-1082>.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. Assessing the Ability of LSTMs to Learn Syntax-Sensitive Dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535, 2016a.

- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535, 2016b.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 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, June 2019a. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/N19-1112>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019b.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- Shiqing Ma, Yingqi Liu, Wen Chuan Lee, Xiangyu Zhang, and Ananth Grama. Mode: Automated neural network model debugging via state differential analysis and input selection. In Alessandro Garci, Corina S. Pasareanu, and Gary T. Leavens, editors, *ESEC/FSE 2018 - Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ESEC/FSE 2018 - Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, pages 175–186. Association for Computing Machinery, Inc, October 2018. doi: 10.1145/3236024.3236082. Publisher Copyright: © 2018 Association for Computing Machinery.; 26th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2018 ; Conference date: 04-11-2018 Through 09-11-2018.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 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>.
- Rebecca Marvin and Tal Linzen. Targeted syntactic evaluation of language models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1151. URL <https://www.aclweb.org/anthology/D18-1151>.
- Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction, 2020.

- Richard Meyes, Constantin Waubert de Puiseau, Andres Posada-Moreno, and Tobias Meisen. Under the hood of neural networks: Characterizing learned representations by functional neuron populations and network ablations. *CoRR*, abs/2004.01254, 2020.
- Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? *CoRR*, abs/1905.10650, 2019. URL <http://arxiv.org/abs/1905.10650>.
- Basel Mousi, Nadir Durrani, and Fahim Dalvi. Can llms facilitate interpretation of pre-trained language models? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Singapore, dec 2023. Association for Computational Linguistics.
- Jesse Mu and Jacob Andreas. Compositional explanations of neurons. *CoRR*, abs/2006.14032, 2020. URL <https://arxiv.org/abs/2006.14032>.
- Seil Na, Yo Joong Choe, Dong-Hyun Lee, and Gunhee Kim. Discovery of natural language concepts in individual units of cnns. *CoRR*, abs/1902.07249, 2019. URL <http://arxiv.org/abs/1902.07249>.
- Dat Quoc Nguyen, Dai Quoc Nguyen, Dang Duc Pham, and Son Bao Pham. Rdrpostagger: A ripple down rules-based part-of-speech tagger. In *Proceedings of the Demonstrations at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 17–20, Gothenburg, Sweden, April 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/E14-2005>.
- Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander Mordvintsev. The building blocks of interpretability. *Distill*, 14(14), 2018. doi: 10.23915/distill.00010. <https://distill.pub/2018/building-blocks>.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. A universal part-of-speech tagset. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2089–2096, Istanbul, Turkey, May 2012. European Language Resources Association (ELRA). URL http://www.lrec-conf.org/proceedings/lrec2012/pdf/274_Paper.pdf.
- Nina Poerner, Benjamin Roth, and Hinrich Schütze. Interpretable textual neuron representations for NLP. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 325–327, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5437. URL <https://www.aclweb.org/anthology/W18-5437>.
- Sai Prasanna, Anna Rogers, and Anna Rumshisky. When BERT Plays the Lottery, All Tickets Are Winning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3208–3229, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.259. URL <https://aclanthology.org/2020.emnlp-main.259>.
- Peng Qian, Xipeng Qiu, and Xuanjing Huang. Analyzing Linguistic Knowledge in Sequential Model of Sentence. In *Proceedings of the 2016 Conference on Empirical Methods in*

- Natural Language Processing*, pages 826–835, Austin, Texas, November 2016a. Association for Computational Linguistics. URL <https://aclweb.org/anthology/D16-1079>.
- Peng Qian, Xipeng Qiu, and Xuanjing Huang. Investigating Language Universal and Specific Properties in Word Embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1478–1488, Berlin, Germany, August 2016b. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P16-1140>.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1264. URL <https://www.aclweb.org/anthology/D16-1264>.
- Marius Rangamani, Akshay; Lindegaard, Tomer Galanti, and Tomaso Poggio. Feature learning in deep classifiers through intermediate neural collapse. 2023.
- Nils Rethmeier, Vageesh Kumar Saxena, and Isabelle Augenstein. Tx-ray: Quantifying and explaining model-knowledge transfer in (un-)supervised NLP. In Ryan P. Adams and Vibhav Gogate, editors, *Proceedings of the Thirty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI 2020, virtual online, August 3-6, 2020*, page 197. AUAI Press, 2020. URL http://www.auai.org/uai2020/proceedings/197_main_paper.pdf.
- Hassan Sajjad, Fahim Dalvi, Nadir Durrani, and Preslav Nakov. Poor man’s BERT: smaller and faster transformer models. *CoRR*, abs/2004.03844, 2020. URL <https://arxiv.org/abs/2004.03844>.
- Hassan Sajjad, Nadir Durrani, Fahim Dalvi, Firoj Alam, Abdul Khan, and Jia Xu. Analyzing encoded concepts in transformer language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3082–3101, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.225. URL <https://aclanthology.org/2022.naacl-main.225>.
- Xing Shi, Inkit Padhi, and Kevin Knight. Does string-based neural MT learn source syntax? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP ’16*, pages 1526–1534, Austin, TX, USA, 2016.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps, 2014.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/D13-1170>.

- Karolina Stańczak, Lucas Torroba Hennigen, Adina Williams, Ryan Cotterell, and Isabelle Augenstein. A latent-variable model for intrinsic probing, 2023.
- Xavier Suau, Luca Zappella, and Nicholas Apostoloff. Finding experts in transformer models. *CoRR*, abs/2005.07647, 2020. URL <https://arxiv.org/abs/2005.07647>.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks, 2017.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. BERT rediscovers the classical NLP pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1452. URL <https://www.aclweb.org/anthology/P19-1452>.
- Erik F. Tjong Kim Sang and Sabine Buchholz. Introduction to the CoNLL-2000 shared task chunking. In *Fourth Conference on Computational Natural Language Learning and the Second Learning Language in Logic Workshop*, 2000. URL <https://www.aclweb.org/anthology/W00-0726>.
- Lucas Torroba Hennigen, Adina Williams, and Ryan Cotterell. Intrinsic probing through dimension selection. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 197–216, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.15. URL <https://www.aclweb.org/anthology/2020.emnlp-main.15>.
- Michael Tsang, Sirisha Rambhatla, and Yan Liu. How does this interaction affect me? interpretable attribution for feature interactions, 2020.
- Mehrdad Valipour, En-Shiun Annie Lee, Jaime R. Jamacaro, and Carolina Bessega. Un-supervised transfer learning via BERT neuron selection. *CoRR*, abs/1912.05308, 2019. URL <http://arxiv.org/abs/1912.05308>.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 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, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1580. URL <https://www.aclweb.org/anthology/P19-1580>.
- Ekaterina Vylomova, Trevor Cohn, Xuanli He, and Gholamreza Haffari. Word Representation Models for Morphologically Rich Languages in Neural Machine Translation. *arXiv preprint arXiv:1606.04217*, 14(14), 2016.
- Thiemo Wambsganss, Christiane Engel, and Hansjörg Fromm. Improving explainability and accuracy through feature engineering: A taxonomy of features in nlp-based machine learning. In *International Conference on Interaction Sciences*, 2021.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 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, November 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5446. URL <https://www.aclweb.org/anthology/W18-5446>.

Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1101. URL <https://www.aclweb.org/anthology/N18-1101>.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. Xlnet: Generalized autoregressive pretraining for language understanding. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 5754–5764, 2019. URL <http://papers.nips.cc/paper/8812-xlnet-generalized-autoregressive-pretraining-for-language-understanding>.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67:301–320, 2005.

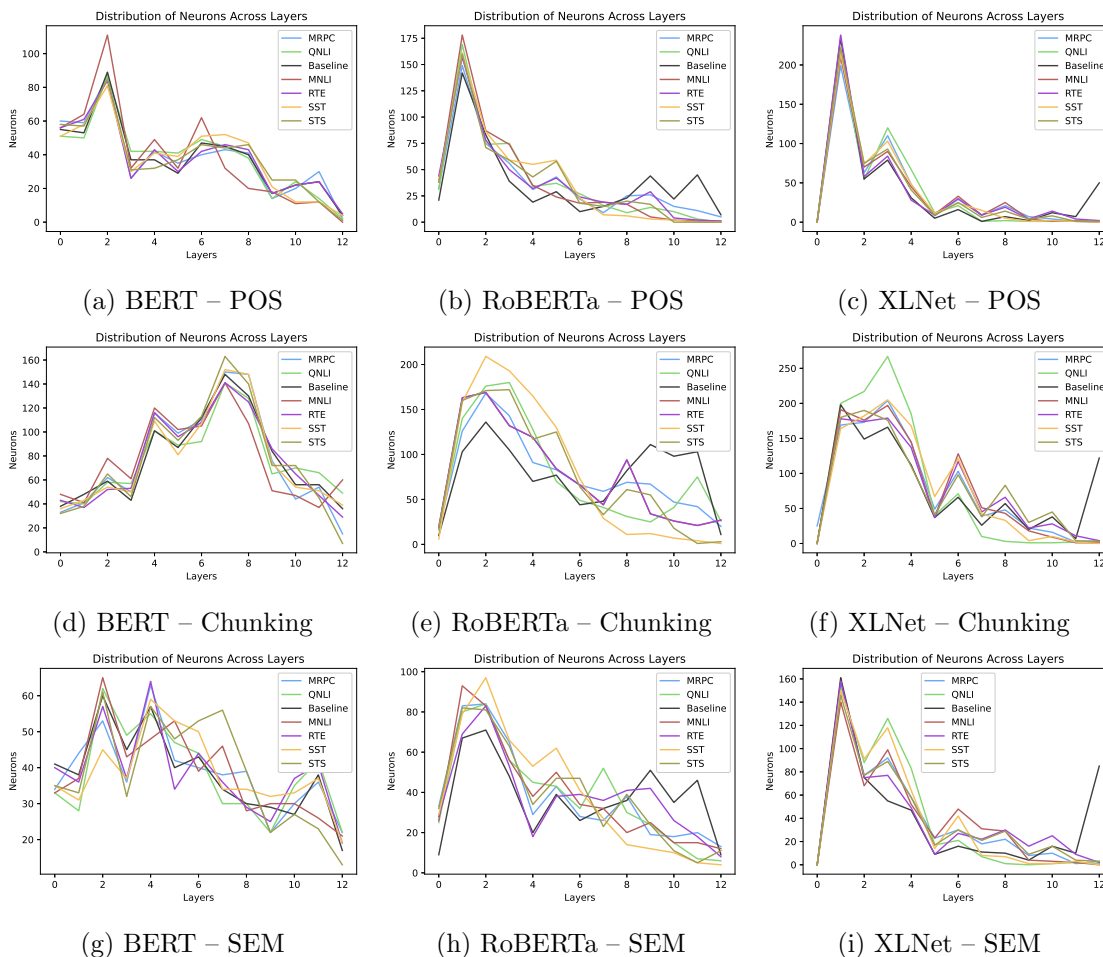


Figure 12: Distribution of Top Neurons across Layers.

Appendix

Appendix A: Data and Representations

We used standard splits for training, development and test data for the 4 linguistic tasks (POS, SEM, Chunking and CCG super tagging) that we used to carry out our analysis on. The splits to preprocess the data are available through git repository⁸ released with Liu et al. (2019a). See Table 8 for statistics. We obtained the understudied pre-trained models from the authors of the paper, through personal communication.

Appendix B: Spread of Neurons

In section 5.1 we presented how distributed or localized neurons are across different properties. We showed that some properties, including closed class words are captured in fewer neurons where as open class words or words that occur in many varying contexts require

8. <https://github.com/nelson-liu/contextual-repr-analysis>

Table 8: Data statistics (number of sentences) on training, development and test sets using in the experiments and the number of tags to be predicted.

Task	Train	Dev	Test	Tags
Suffix	40000	5000	5000	58
POS	36557	1802	1963	44
SEM	36928	5301	10600	73
Chunking	8881	1843	2011	22
CCG	39101	1908	2404	1272

Table 9: Data statistics (number of sentences) on training, development and test sets using in the experiments and the number of POS tags and Syntactic Dependency Relations to be predicted in multilingual experiments.

Task	Train	Dev	Test	Tags
POS (en)	14498	3000	8172	44
POS (de)	14498	3000	8172	52
POS (fr)	11495	3000	3003	13
Syntactic Dependency (en)	11663	1914	3828	49
Syntactic Dependency (de)	14118	1775	1776	35
Syntactic Dependency (fr)	14552	1895	1894	40

more neurons. See Figure 13 for all the properties. In section 5.5 we presented if properties within a task behave similarly across different languages and architectures. In Figure 14 we present results on XLM-RoBERTa where we see a similar evidence as seen in the main results.

Appendix C: Transfer Learning

Section 5.4 presented neuron-wise probing results for selected GLUE tasks. We see a similar pattern across architectures as in Figure 12. As the model is fine-tuned towards downstream, number of salient neurons towards a linguistic property, in the lower layers increase.

DISCOVERING SALIENT NEURONS IN DEEP NLP MODELS

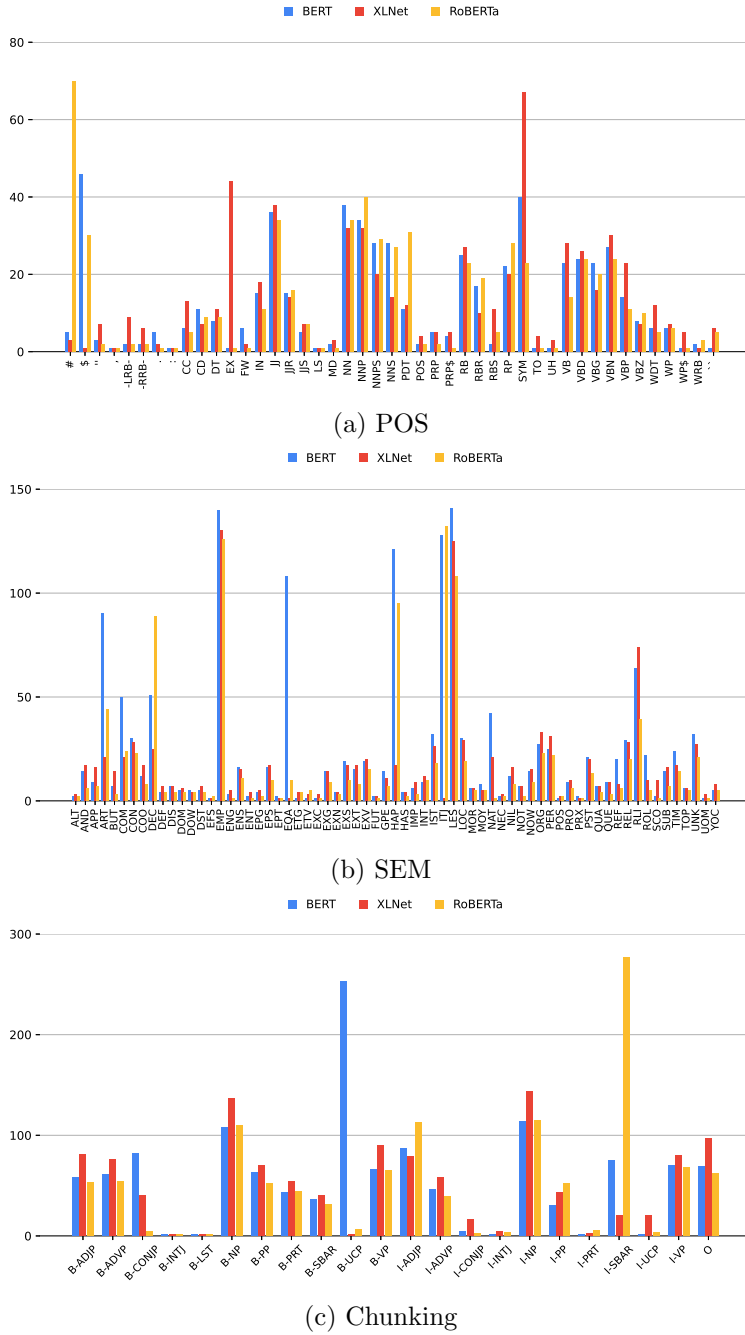
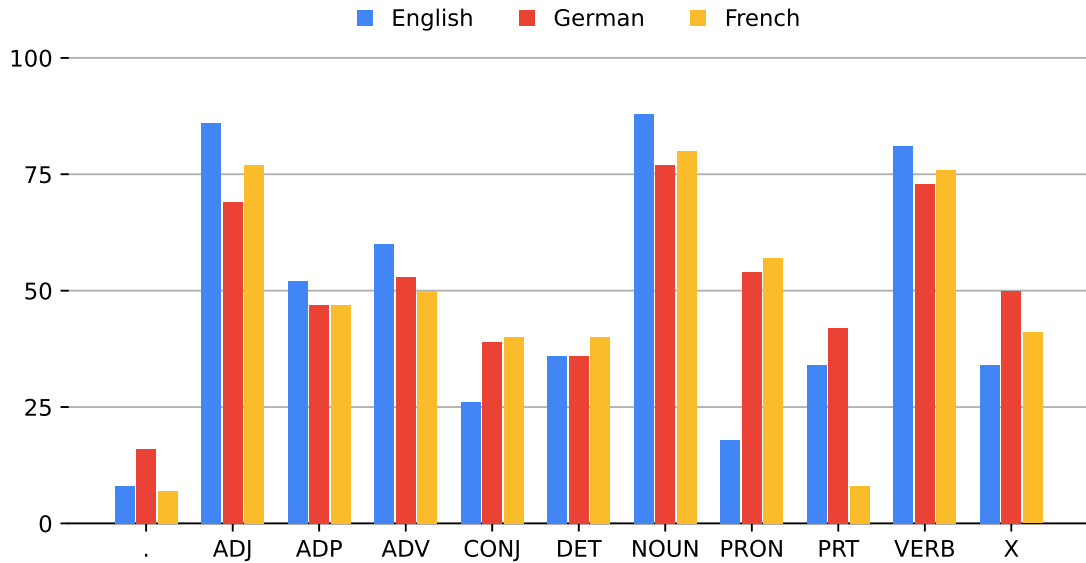
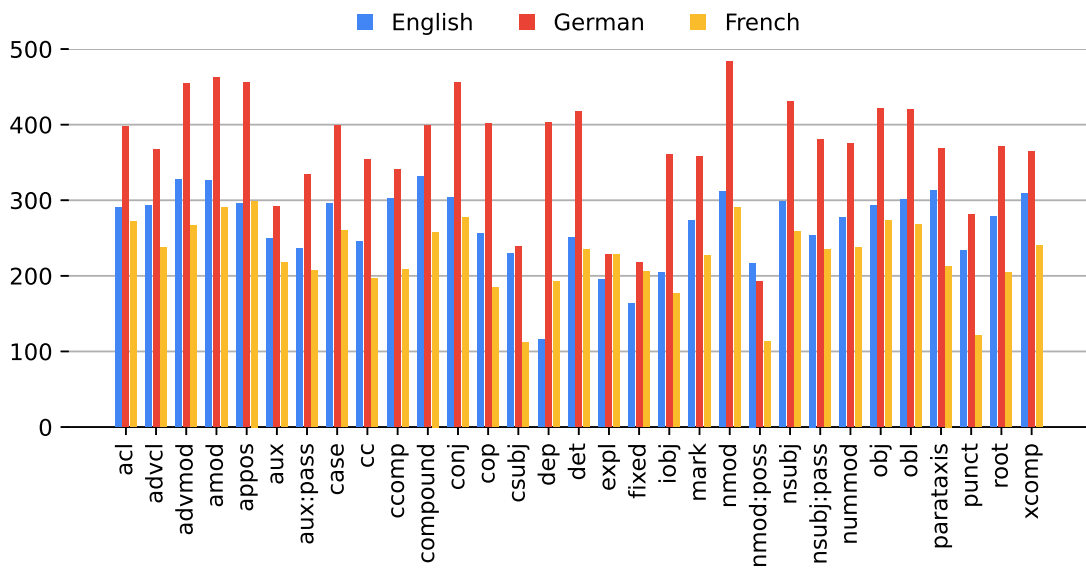


Figure 13: Distribution of Top Neurons Properties.



(a) XLM-RoBERTa – POS tagging



(b) XLM-RoBERTa – Syntactic Dependency Labeling

Figure 14: Distribution of Top Neurons Across POS (a) and Syntactic Dependency Properties in XLM-Roberta.