

The Heads Hypothesis: A Unifying Statistical Approach Towards Understanding Multi-Headed Attention in BERT

Madhura Pande, Aakriti Budhraj, Preksha Nema, Pratyush Kumar and Mitesh M. Khapra

Department of Computer Science and Engineering, IIT Madras, India
 Robert Bosch Center for Data Science and Artificial Intelligence (RBC-DSAI) IIT Madras, India
 {mpande, abudhra, preksha, pratyush, miteshk}@cse.iitm.ac.in

Abstract

Multi-headed attention heads are a mainstay in transformer-based models. Different methods have been proposed to classify the *role* of each attention head based on the relations between tokens which have high pair-wise attention. These roles include syntactic (tokens with some syntactic relation), local (nearby tokens), block (tokens in the same sentence) and delimiter (the special [CLS], [SEP] tokens). There are two main challenges with existing methods for classification: (a) there are no standard scores across studies or across functional roles, and (b) these scores are often average quantities measured across sentences without capturing statistical significance. In this work, we formalize a simple yet effective score that generalizes to all the roles of attention heads and employs hypothesis testing on this score for robust inference. This provides us the right lens to systematically analyze attention heads and confidently comment on many commonly posed questions on analyzing the BERT model. In particular, we comment on the co-location of multiple functional roles in the same attention head, the distribution of attention heads across layers, and effect of fine-tuning for specific NLP tasks on these functional roles. Code is made publicly available at <https://github.com/iitmnlp/heads-hypothesis>

Introduction

In the short span of two years, the BERT model (Devlin et al. 2019) has become a top contender in many NLP tasks. This success has led to a sub-field of research that tries to analyze why BERT works. This sub-field has become so prominent that it has even received its own name - BERTology. The state of BERTology is well captured by the summary in a recent survey paper: “*while BERTology has come a long way it is fair to say we still have more questions than answers about how BERT works*” (Rogers, Kovaleva, and Rumshisky 2020). We see these questions as being broadly classified into three types: (a) what does BERT know, (b) what are BERT outputs sensitive to, and (c) what do the attention heads in BERT attend to? For identifying what BERT knows, a typical strategy is to run probing experiments, for instance by evaluating POS tagging with contextual embeddings from a specific layer of BERT (Liu et al. 2019). For understanding how BERT outputs vary, ablation studies such as

pruning attention heads and dropping layers are performed (Budhraj et al. 2020; Michel, Levy, and Neubig 2019; Sajjad et al. 2020). Quite separate from these two questions, the third question aims to analyze the *intrinsic* properties of the attention heads, which are characteristic of Transformer-based models (Vaswani et al. 2017). Analysis and interpretation of attention patterns has been a matter of debate (Jain and Wallace 2019; Wiegrefe and Pinter 2019). However, it is agreed that studying a single attention head at a time enables us to structurally localize linguistic knowledge within the model. In this work, we survey the current approaches to analyze attention heads, identify some deficiencies, and propose a unifying and statistically robust alternative.

Multi-headed self-attention is characterised by attention weights which specify the weights of other tokens when computing the representation of the current token. A natural curiosity is if these attention weights encode patterns that are linguistically or structurally meaningful. More specifically, analysis aims to classify the role of an attention head into functional roles such as *local* (attending to tokens in a small neighbourhood of the current token), *syntactic* (attending to tokens which are syntactically related to the current token), *delimiter* (attending to delimiter tokens such as [CLS] and [SEP]) and *block* (attending to tokens within the same sentence). Several methods have been proposed for such classification of functional roles of attention heads. However, these methods vary significantly from one work to the other and also amongst the different functional roles. For example, (Kovaleva et al. 2019) use the attention heatmap as an image and a CNN-based classifier to determine whether the image exhibits a local attention pattern. Contrast this with the method used in (Voita et al. 2019) which classifies a head as local if 90% of the attention of the head lies within a small neighbourhood of the current token. Contrast this further with the method used in the same work to identify a syntactic role: A head is labeled as syntactic if the classification accuracy of this head on certain syntactic roles is higher than baseline classifiers. The “higher than” assertion is itself based on the subjectively chosen parameter of 10%.

Thus, existing studies present a diverse set of methods for analysis with their respective parameters which impedes systematic comparison and evolution of ideas. Furthermore, existing studies have an additional challenge: statistical significance. Clearly, as we change the input sequence (one or

more sentences) to the model, the attention patterns across heads would significantly vary. Thus, any claim on the functional role of an attention head must be shown to be consistent across the variation of input sequences. Such analysis is well formalized in the domain of statistics where hypothesis tests must be validated with high confidence. However, no existing work formally presents such analysis for classification of attention heads. This observation is in line with the general call for more rigorous statistical validation across NLP (Azer et al. 2020).

To address these challenges with existing analysis of attention heads, we propose a single unifying metric which is then validated with hypothesis testing to determine functional roles of the attention heads. The metric we propose is quite straightforward and is based on two observations. First, for every hypothesized functional role of attention heads, for a given input sequence, we can identify a target set of tokens which the current token is supposed to attend to. We call this target set of tokens, the attention *sieve*. Second, we need to score the preferential bias of a head in attending to tokens in a sieve. We can aggregate the attention weights of a head to all tokens in the sieve, but we must normalize this against the number of tokens in the sieve and the total number of tokens in the input sequence. This normalized aggregate attention paid by an attention head to the sieve for each of the functional roles is referred to as the *sieve bias*. Then, classifying an attention head to one of the functional roles can be formulated as a hypothesis test that the population mean of the corresponding sieve bias across the population of input sequences exceeds a specified threshold. The threshold here specifies how sharp the attention bias should be to classify the functional role of a head. It is still subjective, but is interpretable and consistent across functional roles.

The above formulation is proposed to unify and systematize the analysis of attention heads. Additionally, it provides us a lens to confidently comment on many questions posed during analysis of attention heads. We take up three main questions in this paper.

1. Are functional roles mutually exclusive? Most existing works (Kovaleva et al. 2019) classify attention heads to one of the different functional roles. This precludes the option that multiple functional roles can co-exist in the same head. There are two reasons to consider this possibility. One, functional roles defined in terms of structure (such as local heads) may overlap with syntactic roles (such as attending to `nsubj`) given that these syntactic relations are often local in the sentences. This is a crucial overlap as recent studies have attempted to exclusively use local attention in Transformer-based models (Raganato, Scherrer, and Tiedemann 2020). Two, a single head may perform multiple roles simultaneously - e.g. attending to two different syntactic relations (such as `dobj` and `amod`). In the paper, we quantify, visualise, and comment on such overlap in functional roles.

2. How are functional roles distributed across layers? Existing works have shown that more syntactic attention heads are found in middle layers as opposed to other layers (Jawahar, Sagot, and Seddah 2019; Hewitt and Manning 2019; Goldberg 2019). Our work with hypothesis testing on sieve bias scores provides a fresh perspective on these studies. We

can define fine-grained functional roles such as separating out the individual syntactic relations and delimiter tokens. Then we can quantify the fraction of heads in each layer than are *skilled* to perform each of these fine-grained roles. In this work, we quantify such skills across layers and reaffirm existing findings.

3. What is the effect of fine-tuning on functional roles?

A common workflow of BERT is to fine-tune the model for specific NLP tasks. Earlier studies (Kovaleva et al. 2019) have shown that only the last few layers specialize during fine-tuning. We can analyse such specialization by quantifying the change in the *sieve bias* across layers and functional roles. In this paper, we compute these changes in sieve bias across four different NLP tasks from the GLUE benchmark.

In summary, we propose a unifying approach to analyze attention heads that generalizes across functional roles and is statistically significant. The new metrics and visualization tools enable a fresh commentary on existing findings from BERTology.

Related Work

As mentioned earlier, the field of BERTology deals with three main questions: (a) what does BERT know? (Lin, Tan, and Frank 2019a; Tenney et al. 2019; Ettinger 2020) (b) what are BERT outputs sensitive to? (Ethayarajh 2019; Wiedemann et al. 2019; Mickus et al. 2019) and (c) what do the attention heads in BERT attend to (Kovaleva et al. 2019; Clark et al. 2019; Cui et al. 2019)? Of these, in this work we focus on the last question. In particular, we are interested in understanding the functional role of heads in BERT. In this section, we review existing work along two different axes.

Types of roles identified: Different works have identified a variety of functional roles for attention heads. These are:

local: heads which attend to tokens in a small neighborhood around the input token, typically, previous and/or next tokens (Clark et al. 2019; Kovaleva et al. 2019; Htut et al. 2019). Some works refer to such heads as diagonal heads (Kovaleva et al. 2019) or positional heads (Voita et al. 2019).
syntactic: heads which attend to tokens which are syntactically related to the input token. The syntactic relations are identified using a dependency parser (Voita et al. 2019; Clark et al. 2019; Htut et al. 2019; Kovaleva et al. 2019; Correia, Niculae, and Martins 2019).

vertical/delimiter: heads which attend to the `[SEP]` and `[CLS]` tokens in the input sequence (Clark et al. 2019; Kovaleva et al. 2019). Some works (Lin, Tan, and Frank 2019b; Htut et al. 2019) specifically ignore the attention on `[SEP]` and `[CLS]` as these are artificially introduced tokens and not a part of the input sentence(s).

block: heads which attend to tokens within the same sentence as opposed to tokens in other sentence before (or after) the `[SEP]` token (Kovaleva et al. 2019).

rare word: heads which attend to rare words, *i.e.*, words with a low frequency in the corpus (Voita et al. 2019).

BPE-merging: heads which attend to siblings of the current token resulting from the tokenisation using BPE (Correia, Niculae, and Martins 2019).

interrogative: heads which attend to question marks at the end of a question (Correia, Niculae, and Martins 2019).

In this work, we focus on *local*, *syntactic*, *block*, and *delimiter* heads which are commonly studied in existing works.

Methods used for classifying a head into a functional role: Once the above functional roles are defined, existing works classify heads into these roles using a scoring function and an appropriate threshold. However, across different works there is no uniformity in the scoring function and the thresholds used for different functional roles. For example, Voita et al. classify a head as local if in at least 90% of the cases, the head pays maximum attention to tokens which are immediately before or after the input token. In contrast, for determining whether a head is syntactic or not, its performance in classifying syntactic roles is compared with a simple position based baseline classifier. A head is then labelled as syntactic if its performance is at least 10% higher than the baseline classifier. Clark et al. also follow a similar strategy wherein they classify a head as syntactic based on its performance on identifying syntactic roles. However, unlike Voita et al. they do not use a threshold of 10% but instead look for heads which substantially outperform a position based baseline classifier. Kovaleva et al. use a CNN based classifier to assign roles to heads based on attention heatmaps. They also propose certain semantic and syntactic tests for attributing roles to heads. For example, they check if certain heads pay more attention to predicates of core frame elements as captured in FrameNet. Existing methods also differ in the treatment of special tokens such as [SEP] and [CLS]. For example, (Htut et al. 2019) and (Lin, Tan, and Frank 2019b) exclude these tokens from their analysis (and thus don’t consider delimiter heads). In contrast, (Clark et al. 2019) show that delimiter heads are prominent in the later layers and thus important. Thus, the mechanisms for determining head’s functional role not only varies from one work to another but also from one functional role to another.

Another limitation with existing work is that they rely on *average* functional scores of heads across a large number of input sequences. Since the attention patterns (e.g., of syntactic heads) vary sensitively from one input sequence to another, this raises the question of statistical significance. Indeed, in our experiments we find that averages are inflated by outliers sufficiently to incorrectly infer functional roles that are not statistically significant. In this work, we address these deficiencies by proposing a statistically sound unified approach which can be adapted to any functional role.

Unifying Analysis of Attention Heads

The BERT Model

We analyze the BERT model which is a stack of layers each consisting of multi-headed self-attention followed by a fully connected network. The self-attention scheme works as follows: For every token in the input sequence an attention head transforms the token’s input embedding into key, value, and query vectors, which are then linearly combined based on attention weights. The outputs of all heads in a layer are concatenated and then passed through the fully connected network. The embeddings for the next layer are generated by combining the output of a layer with a skip connection from the previous layer. We specifically consider the BERT_{BASE}

configuration (Devlin et al. 2019) which has 12 layers with 12 attention heads each, and embedding vector of size 64.

Attention Sieves

As noted, attention heads are classified to have different functional roles: local, syntactic, block, and delimiter. Our primary aim is to define a unifying scheme to analyze attention heads across all functional roles. The first observation is that attention heads for different functional roles are expected to attend to different sets of tokens. We characterize this *set* of tokens for a given attention role as an attention *sieve*. Formally, attention sieve $S_{I,f}(t)$ for input sequence I , functional role f , and token t is the set of all tokens of I that should be attended to compute the embedding of t by an attention head classified to be of functional role f . We illustrate this for the example sequence shown in Figure 1. A local attention head has a clear sieve that does not depend on the content of the input sequence. The local attention sieve for a token are all tokens around its neighbourhood within a distance of say 2, *i.e.*, 5 tokens. For a syntactic attention head we have to first identify all tokens syntactically related to the current token with a dependency parser. These set of tokens would define the syntactic attention sieve. Similarly, we define block and delimiter attention sieves. These attention sieves can be further partitioned to consider sub-types of functional roles. For instance, a delimiter attention sieve can be partitioned to a [CLS] attention sieve and a [SEP] attention sieve. A local attention sieve can be partitioned to a prev-local or a next-local attention sieve that defines locality in only one direction. Thus, for all tokens of an input sequence, we get attention sieves for each functional role.

Sieve Bias Score

For a token t of an input sequence I , an attention head h computes an attention weight $\alpha_h(t, t')$ for each token $t' \in I$ as follows:

$$\alpha_h(t, t') = \text{softmax} \left(\frac{q_t^T k_{t'}}{\sqrt{d_l}} \right),$$

where q_t is the query vector for token t , and $k_{t'}$ is the key vector for token t' in the sentence. d_l is the size of the embedding. This weight captures the relative contribution of t' in computing the next representation of t . If we are to posit that an attention head has a specific functional role, then attention weights to tokens in the sieve of that functional role must be higher. Formally, if attention head h is of functional role f then $\sum_{t_s \in S_{I,f}(t)} \alpha_h(t, t_s)$ should be relatively large. This aggregation of attention weights must be normalized to consider both the number of tokens in a sieve and the number of tokens in the input sequence. For example, the sieve of a local attention role may have 5 tokens, while that of a delimiter role may have only 2. Even within the same attention role, different words may be broken into varying number of tokens leading to different sizes of attention sieves. For attention head h , current token t , input sequence I , and functional role f , we can capture this normalization with a *sieve bias score* as:

$$\beta_h^f(t, I) = \left(\frac{\sum_{t_s \in S_{I,f}(t)} \alpha_h(t, t_s)}{|S_{I,f}(t)|} \right) / \left(\frac{\sum_{t' \in I} \alpha_h(t, t')}{|I|} \right)$$

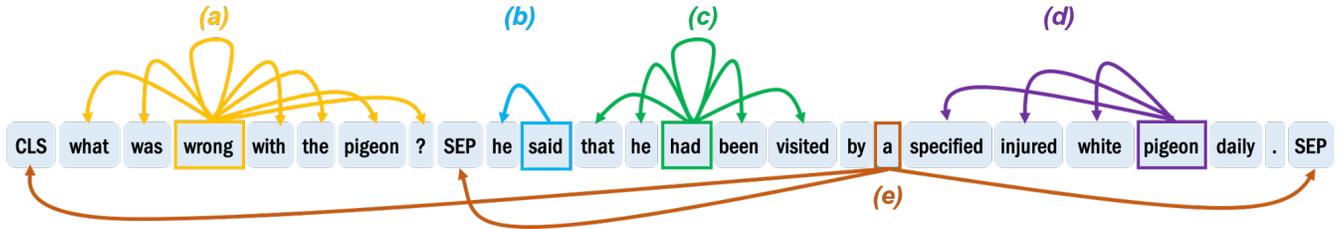


Figure 1: Attention sieves for a sample input sequence: Tokens in boxes show the current token and edges point to other tokens that are in its sieve for various functional roles: (a) block, (b) nsubj (syntactic), (c) local, (d) amod (syntactic) and, (e) delimiter.

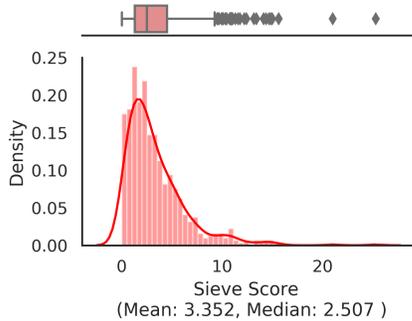


Figure 2: Sieve score distribution across sentences for a head. Note the distribution is skewed with large variation.

In the numerator, we have the average attention weight of the head to tokens within the sieve $t_s \in S_{I,f}(t)$. In the denominator, we have the average attention paid by the head to all tokens in the input sequence. Higher this ratio, higher is the attentive bias to the sieve.

A randomly initialised head would pay roughly equal attention weights to all tokens and thus an sieve bias score of 1 for all functional roles. If upon learning, a head becomes more selective, then its sieve bias score would increase for specific sieves. Our concern then is to identify for each head the functional roles for which sieve bias scores are high.

Hypothesis Testing

We have established a relation between a head’s high sieve bias score and its functional role. But the sieve bias scores $\beta_h^f(t, I)$ vary across input sequences. Can we take the average sieve bias score across many input sequences? This would work if the sieve bias scores are symmetrically distributed with a relatively narrow standard deviation across input sequences. On the contrary, we found that the distributions can be skewed and broad, as shown in Figure 2 for a specific head. In the presence of such variation, simple averaging is error-prone, and we instead need robust statistical inferences. For instance, an assertion based on the average being ≥ 3 would assign the functional role to the head in Figure 2. However, hypothesis testing to assert that population mean is more than 3 would fail. To the best of our knowledge, no current work on analyzing attention heads considers either this variation or a statistical treatment.

We define the null hypothesis as the mean of the sieve bias score β_h^f across all input sequences is less than or equal to some threshold τ .

This threshold τ specifies how sharp the attention bias should be to classify the functional role of a head: Higher the value of τ , stricter is the classification rule. We can evaluate this hypothesis by using a one-tailed test of the mean with z-statistics, since population variance is unknown. To this end, we can measure β_h^f for a large number N of input sequences as the sample. We can then compute the sample mean of β_h^f and derive the p -value using the one-tailed test. As is the convention, we say that the null hypothesis is rejected if the p -value is smaller than 0.05. If the null hypothesis is rejected, it follows that the attention head h is satisfying the functional role f . If the p -value is greater than 0.05, then the test is inconclusive, and we do not ascribe the functional role to the head.

In summary, our proposed method involves the following steps: (a) computing attention sieves for all functional roles, (b) computing sieve bias scores for all pairs of attention heads and functional roles, and (c) performing hypothesis testing on the mean sieve bias score across input sequences to ascribe functional roles to each head.

Results and Discussions

In this section, we first describe our experimental setup and the datasets used for our analysis. We then provide details about the assignment of functional roles to heads using our method described earlier. Finally, based on these assigned roles, we discuss the three questions outlined earlier: i) Are functional roles mutually exclusive? ii) How are functional roles distributed across layers? and iii) What is the effect of fine-tuning on functional roles?

Experimental Setup

We pre-train the BERT_{BASE} model using the English Wikipedia corpus and a subset of of the Project Gutenberg corpus released by (Lahiri 2014). We pre-trained the model for around 300K steps with a sequence length of 128 and another 1K steps with a sequence length of 512 and a batch size of 2K tokens (Devlin et al. 2019) on a single Cloud TPU (v3-8). We then tune this model individually for four NLU tasks from the GLUE benchmark (Wang et al. 2018): *QNLI* (QA Natural Language Inference), *QQP* (paraphrase detection), *MRPC* (paraphrase detection) and *SST-2* (sentiment analy-

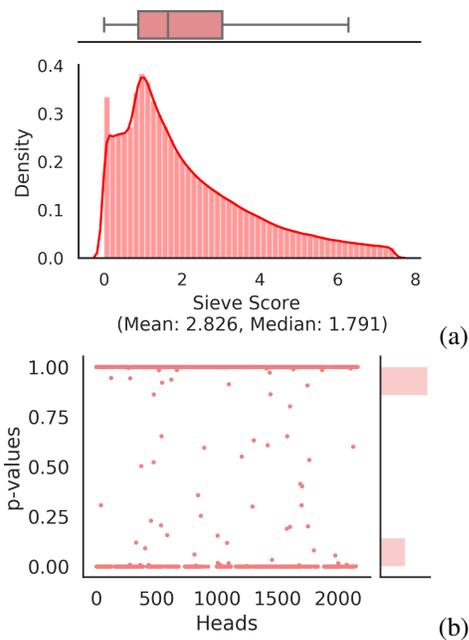


Figure 3: (a) Distribution of sieve scores ($144 \times |\text{tasks}| \times |\text{roles}| \times |\text{num_sents}|$ scores are included) (b) p-values obtained for each head ($144 \times |\text{tasks}| \times |\text{roles}|$) with hypothesis testing. High density of points on $\{0,1\}$ indicate a high confidence level to accept/reject the Null hypothesis.

sis). For each task, we used the standard train and test splits for fine-tuning and evaluation. We used the recommended setting of hyperparameters (Devlin et al. 2019) with batch sizes chosen among $\{32, 128\}$ and learning rates among $\{1e-4, 2e-4\}$. Warm-up was set to 10K steps and LAMB optimizer was used. The performance of the fine-tuned model on these four GLUE tasks is at par with the one reported in Devlin et al. ($BERT_{BASE}$) and suitable for further analysis. We use 1000 input sequences (num_sents) from the standard test sets of the 4 GLUE tasks for all our analysis experiments. We use open-source library spaCy for dependency parsing.

Assigning Functional Roles to Heads

For each task, we feed all the input sequences to the fine-tuned $BERT_{BASE}$ model and compute the sieve bias score $\beta_h^f(t, I)$ for each input sequence I as described earlier. This gives us a sample of 1000 sieve bias scores (one score corresponding to each input sequence). Once we compute these scores for a given head h , and a given functional role f , we evaluate the null hypothesis as discussed in the earlier section based on the obtained p-value for each head. If the null hypothesis is rejected then it means that the head indeed performs this specific functional role. Note that we perform this hypothesis test independently for all heads and the four functional roles. This method thus allows multiple roles to be assigned to the same head.

It is important to discuss the choice of the threshold τ . A large value of τ would lead to the null hypothesis being

rejected for most heads and thus fewer functional roles assigned to heads. A small value of τ would lead to multiple functional roles being assigned to the same head. One approach to set the value of τ is to subjectively decide what constitutes *attentive bias*, i.e., to decide the ratio of attention that a head must pay to tokens in a sieve than to all tokens. For instance, one may suggest that $\tau = 2$ with double the attention to the sieve tokens than to other tokens is a reasonable choice for assigning functional roles to heads. Any such value of τ can be decided and the resulting functional roles can be analysed as we do in the rest of the paper. In addition to this subjective check, we set the value of τ informed by data on the sieve bias scores. Specifically, we plot the cumulative frequency distribution (see Figure 3(a)) of the sieve bias score across all functional roles for 1,000 input sequences, and identify that the mean sieve bias score is around 2.8. We then set τ to be the smallest integer larger than this mean, i.e., 3. We then perform the described hypothesis testing and compute the p-values. We plot the individual p-values across all the input sequences (see Figure 3(b)). We observe that this distribution is distinctly bimodal with values either close to 1 or close to 0. Indeed, there are only 1.8% of points in the range (0.05, 0.95). This indicates that the hypothesis testing made clear inferences: Either the null hypothesis was rejected (due to p-values less than 0.05) or the test was inconclusive but with very high p-values. In the former case, we assign the functional role to the head with high confidence, and in the latter case the head is clearly not satisfying that functional role as indicated by the high p-values. Thus, we set the threshold $\tau = 3$ to meet the two criteria: a subjective criterion of attentive bias and the objective requirement that the hypothesis testing generates clear inferences. Further, note that this threshold is interpretable as it simply means that the attention on the tokens within the sieve is 3 times the attention on tokens outside the sieve. Based on this statistically sound unified approach of classifying heads we now revisit some commonly asked questions while analysing attention heads.

Are Functional Roles Mutually Exclusive?

As mentioned earlier, existing methods typically assign only one functional role to a head. In contrast, as we do the hypothesis test independently for each functional role, the same head can potentially be labeled with multiple functional roles. To find out whether this is really the case, we present a Venn Diagram where heads being labeled with the same role constitute of a set (see Figure 4). We make the following observations from this figure:

Density of delimiter roles: Across tasks, there are significant number of delimiter heads (on average 73.43% across the 4 tasks). Also given their high prevalence, they naturally intersect with other functional heads.

Unskilled heads: There are very few heads which do not perform any functional role (avg. 9.37% across the 4 tasks).

Overlap between local and syntactic roles: There is a significant overlap between local and syntactic heads. More specifically, across the four tasks, 42.1% to 88.8% (with a mean of 70.9%) of syntactic heads overlap with local heads. This observation has an interesting connection to other stud-

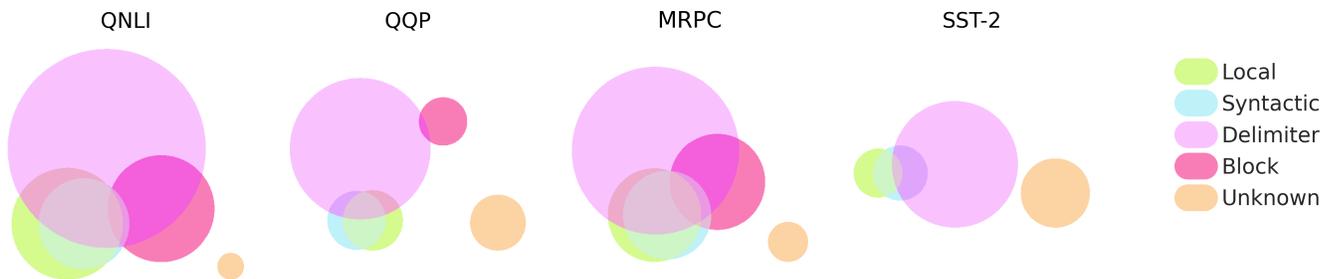


Figure 4: Venn diagram with each circle representing a set of heads with a specific role. The intersection between various sets indicate that many heads are multi-functional. Note that the size of the circle is proportional to the cardinality of the set.

ies which show that restricting attention to local patterns can still lead to good performance (Yang et al. 2019). Our findings lend more credence to this observation, as they indicate that by attending to local tokens a head invariably also attends to syntactically related tokens. This is simply because many syntactically related tokens are in a small neighborhood around the input token. To further strengthen this claim, we compute the Spearman correlation between the local and syntactic sieve scores. We observe a very high correlation score of 0.78, 0.81, 0.85, 0.73 scores for *QNLI*, *QQP*, *MRPC*, *SST-2* tasks respectively, with a p -value of 0. We further split the syntactic and location heads based on type of syntactic relations and location of attended token (previous or next) respectively. We observe that *nsubj* heads have a high overlap with *next-loc* heads and *obj*-heads have a high overlap with the *prev-loc* heads.

Role of block heads: There is only a marginal overlap between block heads and syntactic/local heads indicating that these specialisations are unique. Note that there are no block heads in the *SST-2* task as it is a single sequence task.

How Are Functional Roles Distributed across Layers?

Existing works have shown that the middle layers in BERT are more important as they contain many syntactic heads (Hewitt and Manning 2019; Goldberg 2019; Clark et al. 2019). Our results suggest that while this is largely true, there is significant variance across tasks. To understand this, we refer the reader to Figure 5. Every larger gray colored cell in this figure corresponds to one head and each row corresponds to one layer. Thus, there are a total of $12 \times 12 = 144$ larger gray squares in each plot. We further divide each of these larger gray squares into 9 smaller squares. Each of these 9 squares corresponds to a different functional role represented by a different color as defined in the legend. For this analysis, we have split the four coarse functional roles into more fine-grained functional roles. In particular, we have split the delimiter role into [SEP] and [CLS] roles. Similarly, we have split the syntactic role into *nsubj*, *obj*, *amod* and *advmod* roles. Squares which are predominantly gray correspond to heads which perform very few functional roles. On the other hand, multi-colored squares correspond to *multi-skilled* heads, i.e., heads performing multiple func-

tional roles. For example, in the first plot in Figure 5 (*QNLI*), the first head in *Layer7* performs eight different functional roles, whereas the last head in the same layer performs only one functional role [SEP]. Also, to aid better visualisation, in each layer we have sorted the heads based on the number of functional roles they perform (hence, the more multi-colored heads appear at the beginning of the row). Based on these plots we make the following observations:

Multi-skilled heads across layers: The number of *multi-skilled* heads is higher in the middle layers (layers 5 to 9) for *MRPC*, *QNLI* and *SST-2* tasks. However, for *QQP* the number of multi-skilled heads is relatively higher in the initial layers. This suggests that the role of heads across layers varies from one task to another. Further, these observations may differ since earlier studies have been mostly based on analyzing a single functional role at a time.

Distribution of delimiter heads: Across tasks we find a large number of delimiter heads ([SEP] and/or [CLS]) across all layers. In fact, such delimiter heads are often the only functional heads in the last three layers across different tasks. Interestingly, the first layer contains very few (if any) delimiter heads. This may be because the delimiter tokens [CLS] and [SEP] start accumulating meaningful information from the entire sequence only after the first layer.

Distribution of block heads. For all the tasks which take two sentences as inputs (*MRPC*, *QNLI* and *QQP*) block heads are prevalent in the zeroth layer. However, their presence in the later layers varies from one task to another. For example, in both *QNLI* and *MRPC*, block heads consistently show up in layers 1 to 4, in addition to layer 0. However, in *QQP* block heads do not show up in the middle layers but are prevalent in the first and last layer. Lastly, for *MRPC* block heads consistently show up across all layers.

Distribution of syntactic heads. Across all tasks there are at least 3 syntactic heads (*nsubj*, *obj*, *amod* and *advmod*) for all layers (with very few exceptions). Across tasks, the prevalence of syntactic heads is higher in layers 6 to 8 which is in line with observations made in other studies (Rogers, Kovaleva, and Rumshisky 2020; Jawahar, Sagot, and Seddah 2019; Liu et al. 2019). However, our findings suggest that the heads in these layers do not just perform syntactic roles but are highly *multi-skilled*. Lastly, across tasks the last two layers have very few syntactic or multi-functional heads.

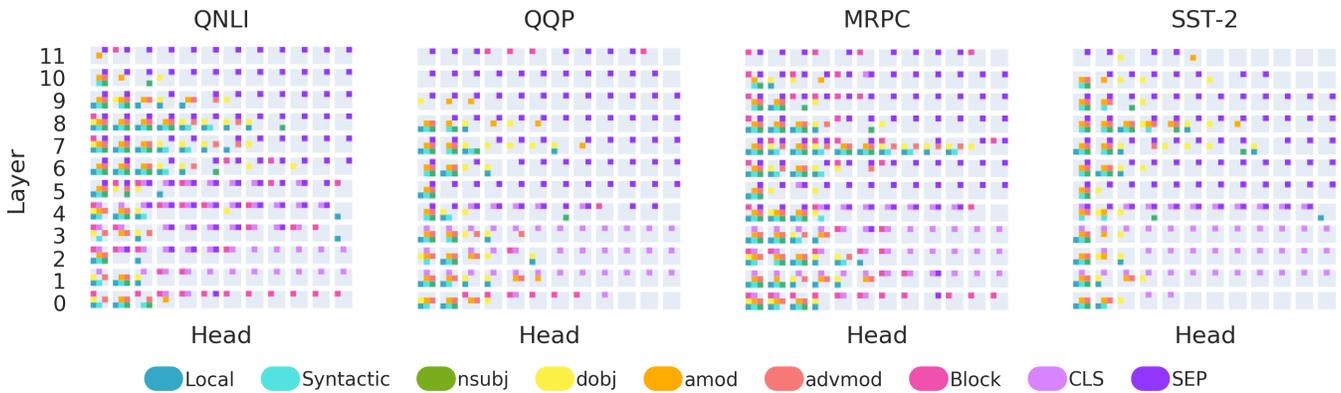


Figure 5: Visualization of the functional roles of the attention heads of the BERT model fine-tuned on four different tasks. In each layer, the heads are arranged in a sorted order (decreasing) of their degree of multi-functionality.

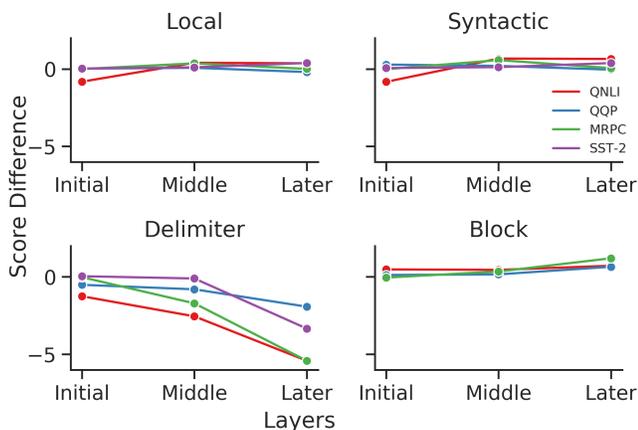


Figure 6: Difference between the average sieve bias score after and before fine-tuning across functional roles and tasks.

What is the Effect of Fine-Tuning on Functional Roles?

Existing works have shown that the final layers of BERT are more task specific. For example, Kovaleva et al. show that the weights in the final layers change a lot during fine-tuning, whereas, Hao et al. show that the performance is less sensitive to weights in the middle layers. To substantiate these findings we measure the change in the average sieve bias scores for the 4 functional roles across initial, middle, and final sets of 4 layers for the 4 tasks (see Figure 6). A positive (negative) value indicates that the average sieve bias score increases (decreases) after fine-tuning. We make the following observations from these plots:

Functional roles get redistributed in the final layers: We observe that in the final layers, there is a decrease in the average delimiter bias score (*i.e.*, avg. $\beta_h^{delimiter}$ across all heads in these layers). On further investigation, we find that this decrease is mainly due to a lower attention on the [SEP] but the attention on the [CLS] token does not change much. This seems to substantiate the hypothesis made in Clark

et al. that the [SEP] tokens serve as a no-op indicator and hence the attentive bias to those tokens is reducing as fine-tuning specializes the final layers. We also find that there is an increase in the average block bias score (*i.e.*, avg. β_h^{block} across all heads in these layers). Thus, the later layers are specialising to generally attend to tokens within their sentence without specific syntactic bias.

Functional roles do not change in the initial and middle layers: In contrast to the final layers, we observe that average bias scores do not change much in the initial layers that is in line with observations made in other studies. Hao et al. (2019) observe low sensitivity to the weights in middle layers. However, we offer a different perspective that bias scores in initial and middle layers do not change much.

Conclusion

In this work, we present a unified and formal approach for analysing the functional roles of attention heads. In particular, the sieve bias score generalizes across functional roles and hypothesis tests ensure statistical significance. This systematises the study of attention heads in BERTology. Based on this analysis, we make several observations. Delimiter heads are prevalent in the network and thus overlap with other functional heads. Syntactic heads often are also local heads which explains claims made in other studies about the sufficiency of local attention patterns (Yang et al. 2018, 2019; Wu et al. 2019). Middle layers have most number of multi-skilled heads, in line with observations made in other studies about the drop in performance while pruning heads in middle layers. Lastly, we are able to carefully study the impact of task-specific fine-tuning on the roles of the attention heads. We find that later layers have the largest change in their roles and specifically the number of delimiter heads drops sharply indicating specialization for task-specific roles. This again supports existing claims about the importance of later layers for specific tasks while providing a fresh perspective based on functional roles. In summary, our formal unified approach to analysing attention heads provides us the right lens to confidently comment on various questions about the functional roles of attention heads.

Acknowledgements

We thank Google for providing us with free TPUs (through TFRC program) which tirelessly ran our code day in and day out. We also thank them for supporting Preksha Nema through their Ph.D. Fellowship program. We are grateful to the Department of Computer Science, RISE Lab, Robert Bosch Center for Data Sciences and Artificial Intelligence (RBC-DSAI), and IIT Madras. We also thank the anonymous reviewers for their valuable suggestions and Madhura Pande acknowledges Maneesha for her support.

References

- Azer, E. S.; Khashabi, D.; Sabharwal, A.; and Roth, D. 2020. Not All Claims are Created Equal: Choosing the Right Statistical Approach to Assess Hypotheses. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5715–5725.
- Budhreja, A.; Pande, M.; Nema, P.; Kumar, P.; and Khapra, M. M. 2020. On the Weak Link between Importance and Prunability of Attention Heads. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3230–3235.
- Clark, K.; Khandelwal, U.; Levy, O.; and Manning, C. D. 2019. What Does BERT Look At? An Analysis of BERT’s Attention. *arXiv preprint arXiv:1906.04341* .
- Correia, G. M.; Niculae, V.; and Martins, A. F. 2019. Adaptively sparse transformers. *arXiv preprint arXiv:1909.00015* .
- Cui, B.; Li, Y.; Chen, M.; and Zhang, Z. 2019. Fine-tune BERT with Sparse Self-Attention Mechanism. 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)*, 3539–3544.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. doi:10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- Ethayarajh, K. 2019. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. *arXiv preprint arXiv:1909.00512* .
- Ettinger, A. 2020. What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics* 8: 34–48.
- Goldberg, Y. 2019. Assessing BERT’s Syntactic Abilities. *ArXiv abs/1901.05287*.
- Hao, Y.; Dong, L.; Wei, F.; and Xu, K. 2019. Visualizing and Understanding the Effectiveness of BERT. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 4141–4150. Association for Computational Linguistics. doi:10.18653/v1/D19-1424. URL <https://doi.org/10.18653/v1/D19-1424>.
- Hewitt, J.; and Manning, C. D. 2019. A Structural Probe for Finding Syntax in Word Representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4129–4138. Minneapolis, Minnesota: Association for Computational Linguistics. doi:10.18653/v1/N19-1419. URL <https://www.aclweb.org/anthology/N19-1419>.
- Htut, P. M.; Phang, J.; Bordia, S.; and Bowman, S. R. 2019. Do Attention Heads in BERT Track Syntactic Dependencies? *CoRR abs/1911.12246*. URL <http://arxiv.org/abs/1911.12246>.
- Jain, S.; and Wallace, B. C. 2019. Attention is not explanation. *arXiv preprint arXiv:1902.10186* .
- Jawahar, G.; Sagot, B.; and Seddah, D. 2019. What Does BERT Learn about the Structure of Language? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3651–3657. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/P19-1356. URL <https://www.aclweb.org/anthology/P19-1356>.
- Kovaleva, O.; Romanov, A.; Rogers, A.; and Rumshisky, A. 2019. Revealing the dark secrets of BERT. *arXiv preprint arXiv:1908.08593* .
- Lahiri, S. 2014. Complexity of Word Collocation Networks: A Preliminary Structural Analysis. In *Proceedings of the Student Research Workshop at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 96–105. Gothenburg, Sweden: Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/E14-3011>.
- Lin, Y.; Tan, Y. C.; and Frank, R. 2019a. Open Sesame: Getting Inside BERT’s Linguistic Knowledge. *arXiv preprint arXiv:1906.01698* .
- Lin, Y.; Tan, Y. C.; and Frank, R. 2019b. Open Sesame: Getting Inside BERT’s Linguistic Knowledge. *CoRR abs/1906.01698*. URL <http://arxiv.org/abs/1906.01698>.
- Liu, N. F.; Gardner, M.; Belinkov, Y.; Peters, M. E.; and Smith, N. A. 2019. Linguistic knowledge and transferability of contextual representations. *arXiv preprint arXiv:1903.08855* .
- Michel, P.; Levy, O.; and Neubig, G. 2019. Are Sixteen Heads Really Better than One? In *Advances in Neural Information Processing Systems*, 14014–14024.
- Mickus, T.; Paperno, D.; Constant, M.; and van Deemter, K. 2019. What do you mean, BERT? Assessing BERT as a Distributional Semantics Model. *arXiv preprint arXiv:1911.05758* .

- Raganato, A.; Scherrer, Y.; and Tiedemann, J. 2020. Fixed encoder self-attention patterns in transformer-based machine translation. *arXiv preprint arXiv:2002.10260* .
- Rogers, A.; Kovaleva, O.; and Rumshisky, A. 2020. A Primer in BERTology: What we know about how BERT works. *ArXiv abs/2002.12327*.
- Sajjad, H.; Dalvi, F.; Durrani, N.; and Nakov, P. 2020. Poor Man’s BERT: Smaller and Faster Transformer Models. *arXiv preprint arXiv:2004.03844* .
- Tenney, I.; Xia, P.; Chen, B.; Wang, A.; Poliak, A.; McCoy, R. T.; Kim, N.; Van Durme, B.; Bowman, S. R.; Das, D.; et al. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316* .
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Voita, E.; Talbot, D.; Moiseev, F.; Sennrich, R.; and Titov, I. 2019. Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned. *arXiv preprint arXiv:1905.09418* .
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461* .
- Wiedemann, G.; Remus, S.; Chawla, A.; and Biemann, C. 2019. Does BERT make any sense? Interpretable word sense disambiguation with contextualized embeddings. *arXiv preprint arXiv:1909.10430* .
- Wiegrefe, S.; and Pinter, Y. 2019. Attention is not not explanation. *arXiv preprint arXiv:1908.04626* .
- Wu, F.; Fan, A.; Baevski, A.; Dauphin, Y. N.; and Auli, M. 2019. Pay less attention with lightweight and dynamic convolutions. *arXiv preprint arXiv:1901.10430* .
- Yang, B.; Tu, Z.; Wong, D. F.; Meng, F.; Chao, L. S.; and Zhang, T. 2018. Modeling Localness for Self-Attention Networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4449–4458. Brussels, Belgium: Association for Computational Linguistics. doi:10.18653/v1/D18-1475. URL <https://www.aclweb.org/anthology/D18-1475>.
- Yang, B.; Wang, L.; Wong, D.; Chao, L. S.; and Tu, Z. 2019. Convolutional self-attention networks. *arXiv preprint arXiv:1904.03107* .