

Derivational Probing: Unveiling the Layer-wise Construction of Syntactic Structures in Neural Language Models

Anonymous ACL submission

Abstract

Recent work has demonstrated that neural language models encode syntactic structures in their internal *representations*, yet the *derivations* by which these structures are built across layers remain poorly understood. In this paper, we introduce *Derivational Probing* to investigate how micro-syntactic structures (e.g., subject phrases) and macro-syntactic structures (e.g., the relationship between the root verb and its direct dependents) are progressively constructed as word embeddings propagate upward across layers. Our experiments on BERT reveal a clear bottom-up derivation: micro-syntactic dependencies emerge in lower layers and are gradually integrated into a coherent macro-syntactic structure in higher layers. Furthermore, an analysis on a subject-verb agreement task shows that the timing of macro-structure formation is critical for performance, suggesting an optimal intermediate range for integrating global syntactic information.

1 Introduction

Neural language models have achieved remarkable success across a wide range of natural language processing tasks. However, significant uncertainty remains regarding what these models truly learn and how they represent linguistic knowledge. This has spurred extensive research aimed at probing the linguistic capabilities of neural language models (Zhao et al., 2024; Chang and Bergen, 2024).

A prominent line of inquiry is *structural probing*, which directly analyzes word embeddings to uncover latent syntactic structures. For example, Hewitt and Manning (2019) demonstrated that the geometric organization of BERT’s word embedding space encodes syntactic distances defined over dependency parse trees, providing evidence that the model captures syntactic information. Yet, such work typically focuses on the static representations of the whole syntactic structures rather than the

dynamic derivations by which these syntactic structures are built across layers. Understanding not just the resulting representations but also how they are built across layers is essential for a more comprehensive understanding and could also lead to better insights into how these representations are used.

Meanwhile, Tenney et al. (2019) introduced the *expected layer* metric and investigated how different layers in BERT encode different types of linguistic information (e.g., part-of-speech tagging, syntactic parsing, semantic role labeling, and coreference resolution), revealing that the model encodes linguistic abstractions in a manner reflecting a traditional NLP pipeline. However, their approach primarily relied on coarse-grained task accuracy measures, capturing only the overall effectiveness of each layer rather than examining the detailed, layer-wise construction of specific syntactic structures. Consequently, how the syntactic structures are built across layers remains underexplored.

In this paper, we fill this gap by introducing *Derivational Probing*—a method that combines structural probing with an expected layer metric to trace the derivation process of syntactic structures in neural language models (Figure 1). Our approach allows us to examine how *micro-syntactic* structures (e.g., subject and object noun phrases, prepositional phrases) and *macro-syntactic* structures (e.g., the relationship between the root verb and its direct dependents) are built progressively across the layers of a model.

Applying Derivational Probing to BERT (Devlin et al., 2019), our experiments reveal a bottom-up derivation strategy, in which micro-syntactic dependencies emerge in lower layers and are gradually integrated into a coherent macro-structure at higher layers. Furthermore, our detailed analysis on a subject-verb agreement task shows that even when the final syntactic structure is correct, the specific layers at which the macro-structure is constructed

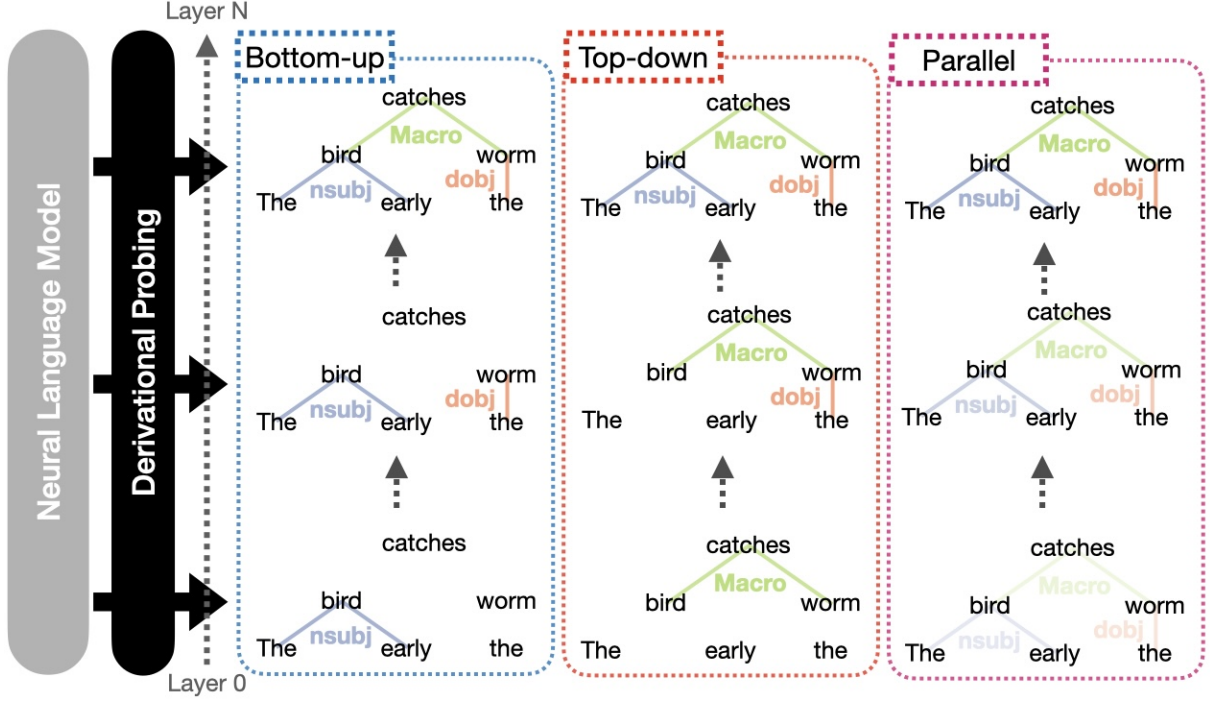


Figure 1: Derivational Probing investigates how syntactic structures are progressively constructed across the layers in neural language models. We illustrate three hypothesized strategies for layer-wise derivation. **Bottom-up:** Micro-syntactic structures, such as subject or prepositional phrases, form in lower layers and the macro-syntactic structure is formed at higher layers. **Top-down:** The macro-syntactic structure is identified first, with micro-syntactic structures refined in subsequent layers. **Parallel:** Micro- and macro-syntactic structures emerge at roughly the same timing.

significantly affect performance. This suggests the existence of an optimal intermediate stage for integrating macro-syntactic information.

Overall, our findings offer new insights into the internal mechanisms by which neural language models construct syntactic trees and underscore the importance of examining syntactic structure formation across layers to improve model interpretability.

2 Related Work

Attention-based analyses (e.g., Clark et al., 2019; Vig and Belinkov, 2019) have demonstrated that certain transformer heads tend to align with dependency relations, providing evidence that Transformer language models capture linguistic dependency relations in their attention weights.

In contrast, Hewitt and Manning (2019) introduced a structural probe that learns a linear transformation from hidden representations into a space where Euclidean distances reflect dependency tree distances. This approach revealed that full syntactic trees are implicitly encoded in models such as BERT. Building on this, later work refined the ap-

proach by incorporating non-linear mappings (e.g., White et al., 2021), enforcing constraints such as orthogonality (Limisiewicz and Mareček, 2021), and using a controlled corpus to isolate the effect of syntax (Maudslay and Cotterell, 2021).

Other studies have refined structural probing by quantifying context-dependent syntactic signals in deeper layers—for example, conditional probing (Hewitt et al., 2021) and information gain metrics (Kunz and Kuhlmann, 2022)—but these methods focus on the performance of specific probing tasks (e.g., POS-tagging) rather than where the syntactic structures are constructed.

In contrast, our proposed method specifically tracks how each subgraph in the syntactic tree develops as information propagates through the network layers. By analyzing the evolution of individual syntactic components—from micro-syntactic structures to the assembly of the macro-syntactic structure—we offer a more granular perspective on the incremental construction of syntax, complementing and extending previous layer-wise analyses.

3 Technical Preliminaries

In this section, we review foundational methods from prior research: structural probing to assess the presence and quality of syntactic representations and the expected layer metric for quantifying how linguistic information gradually builds up across successive layers within language models.

3.1 Structural Probing

Hewitt and Manning (2019) introduced the *structural probe* as a method to evaluate whether contextual word representations encode syntactic information. Given a sentence $s = w_1 \cdots w_t$, each token is represented by a contextual embedding $\mathbf{h}_i \in \mathbb{R}^d$ (e.g., the output embedding of a model like BERT). The goal of the structural probe is to find a linear transformation that maps these embeddings to a space where the Euclidean distances approximate the true syntactic distances between words.

Specifically, for any two words w_i and w_j in a sentence, we define the transformed distance as:

$$d_{\mathbf{B}}(\mathbf{h}_i, \mathbf{h}_j) = \|\mathbf{B}\mathbf{h}_i - \mathbf{B}\mathbf{h}_j\|_2, \quad (1)$$

where $\mathbf{B} \in \mathbb{R}^{d' \times d}$ is a learnable projection matrix. The true syntactic distance, Δ_{ij} , is typically defined as the number of edges on the shortest path between w_i and w_j in the dependency parse tree of the sentence. The probe is trained by minimizing an objective that penalizes the discrepancy between the predicted distances and Δ_{ij} :

$$\mathcal{L} = \frac{1}{|s|^2} \sum_{i=1}^{|s|} \sum_{j=i+1}^{|s|} |\Delta_{ij} - d_{\mathbf{B}}(\mathbf{h}_i, \mathbf{h}_j)|. \quad (2)$$

This formulation encourages the linear transformation \mathbf{B} to capture the syntactic structure encoded in the contextual representations, enabling the recovery of parse trees via Prim’s (1957) algorithm, a greedy algorithm that constructs minimum spanning trees by iteratively adding the lowest-weight edge connecting a new node to the growing tree.

3.2 Expected Layer

The *expected layer* metric introduced by Tenney et al. (2019) was initially developed to identify the layers within BERT responsible for solving various linguistic tasks. Specifically, the metric was used to capture at which layers broad linguistic abilities (e.g., part-of-speech tagging, syntactic parsing, semantic role labeling) emerge.

Formally, Tenney et al. used a *scalar mixing* technique, aggregating information from the bottom layer up to a given layer:

$$\mathbf{m}_i^{(k)} = \gamma \sum_{\ell=0}^k \text{softmax}(\mathbf{a})^{(\ell)} \mathbf{h}_i^{(\ell)}, \quad (3)$$

where $\text{softmax}(\mathbf{a})^{(\ell)} (\mathbf{a} \in \mathbb{R}^{k+1})$ is scalar mixing weights and γ is a learnable scaling factor, following Peters et al. (2018).

By measuring performance improvements across layers, $S(\ell)$, Tenney et al. defined the expected layer to reflect the layer at which the relevant linguistic task information is predominantly captured:

$$E[\ell] = \frac{\sum_{\ell=1}^L \ell \cdot (S(\ell) - S(\ell - 1))}{\sum_{\ell=1}^L (S(\ell) - S(\ell - 1))}. \quad (4)$$

Thus, the expected layer metric was initially proposed to broadly characterize the hierarchical progression of different linguistic capabilities within transformer models, rather than to pinpoint the exact layers at which specific syntactic structures are built.

4 Derivational Probing

Building upon these prior techniques, we propose *Derivational Probing*, a novel method explicitly designed to investigate the dynamic construction of syntactic structures across the layers of neural language models.

Our approach effectively combines expected layer metric (Tenney et al., 2019) with the structural probing (Hewitt and Manning, 2019), enabling a detailed analysis of how syntactic information accumulates across model layers. Specifically, for each layer, we use scalar-mixed embeddings as defined in Eq. (3) and compute pairwise distances:

$$d_{\mathbf{B}_k}(\mathbf{m}_i^{(k)}, \mathbf{m}_j^{(k)}) = \|\mathbf{B}_k \mathbf{m}_i^{(k)} - \mathbf{B}_k \mathbf{m}_j^{(k)}\|_2. \quad (5)$$

We then train the transformation matrix \mathbf{B}_k to minimize discrepancies with true dependency parse distances, analogous to structural probing.

This integration allows us to calculate the expected layer for each syntactic subgraph (micro- and macro-syntactic structures defined in detail later) and perform a fine-grained, quantitative analysis of their construction across model layers. We use the Unlabeled Undirected Attachment Score (UUAS) for each layer ℓ as $S(\ell)$, defined as the proportion of correctly predicted edges to the total

number of edges in the reference dependency parse, without considering edge labels or direction.

To better understand the derivation strategy that models employ when constructing a syntactic tree, we introduce a distinction between *macro-syntactic structures* (the root verb with its direct dependents) and *micro-syntactic structures* (local components, such as subordinate phrases like *nsubj*) (Figure 2). This distinction is motivated by our interest in whether models construct syntactic trees top-down, bottom-up, or in a parallel fashion. To empirically evaluate which of these hypotheses is most plausible, we adopt the following methodological approach: For both micro-syntactic structures and macro-syntactic structures, we (1) construct the full parse tree using a minimum spanning tree algorithm, (2) extract the relevant edges (as highlighted in Figure 2), and (3) compute the UUAS by comparing these edges to the reference parse. By tracking UUAS improvements across layers, we calculate the expected layer $E[l]$ for each structure, revealing the layers at which different syntactic subgraphs are effectively constructed.

We next provide detailed descriptions of each hypothesis.

Bottom-up derivation. A bottom-up derivation first constructs micro-syntactic structures and subsequently integrates these into macro-syntactic structures, ultimately forming a complete dependency tree. We refer to this as a “bottom-up derivation” because it resembles the construction order of the arc-standard transition-based dependency parser (Nivre, 2004). Arc-standard parsing utilizes a stack-based transition system where dependents must be fully processed and attached to their heads before those heads themselves are incorporated into macro-syntactic structures. Under this hypothesis, models initially identify micro-syntactic structures—such as the internal phrase structures of subjects and objects—in lower layers, which are then progressively combined into a coherent macro-syntactic hierarchy at higher layers.

Top-down derivation. A top-down derivation, in contrast, begins by establishing macro-syntactic structures and subsequently refines these by incorporating detailed micro-syntactic dependencies. We term this approach a “top-down derivation” because its construction order aligns closely with the head-driven transition-based parser proposed by Hayashi et al. (2012). Their algorithm explicitly predicts dependent nodes from head nodes, progres-

sively building syntactic structures from head to dependent, thus genuinely following a top-down, predictive parsing order. Under this hypothesis, models prioritize the recognition of macro-syntactic structures before refining micro-syntactic structures.

Parallel derivation. Finally, an alternative hypothesis is that models construct micro- and macro-syntactic structures concurrently, with local dependencies and the global structure forming at roughly the same rate across layers. This hypothesis is less clearly aligned with traditional dependency parsing algorithms, as most classical approaches tend to favor either bottom-up or top-down derivations.

Notes on the term “derivation”. Here, we explicitly use the term *derivation* (strategy) throughout this paper rather than “parsing strategy” to clearly distinguish two related but distinct concepts. While “parsing strategy” generally refers to methodological choices for *incrementally* constructing a parse tree (such as bottom-up or top-down), our use of “derivation” specifically captures an *atemporal* process describing how syntactic structures progressively emerge across the internal layers of a language model given the full sentence context, emphasizing *layer-wise* structural development rather than sequential, *left-to-right* incremental processing.




Figure 2: *Macro-syntactic structure* (“Macro”) and *micro-syntactic structures* (“nsubj” and “dobj”).

5 Experimental Setup

5.1 Data

We utilize the Wikitext-103 dataset (Merity et al., 2016) as our primary source of natural language, parsing each sentence with spaCy’s dependency parser (EN_CORE_WEB_LG) (Honnibal et al., 2020). To focus on the language model’s ability to construct syntactic structures in a clear-cut setting, we restrict our analysis to single-clause sentences by excluding those with relative clauses or clausal

subjects. Additionally, we filter out sentences containing dependency relations such as “dep” (unclassified dependents) and punctuation marks other than sentence-final punctuation to minimize noise.

Following the definitions introduced in the previous section (§4), we group sentences based on dependency relations emanating from the root verb, thereby distinguishing between the overall (macro-syntactic; Macro) structure and subordinate (micro-syntactic) structures. We retain only those groups that represent more than 10% of the data, focusing our analysis on the predominant structure sets. This filtering results in four primary structure sets (See App. A for examples): (1) Macro with micro relations *nsubj* and *dobj*; (2) Macro with micro relations *nsubj* and *prep*; (3) Macro with micro relations *nsubj* and *attr*; and (4) Macro with micro relations *nsubj*, *prep*, and *dobj*.

From the resulting dataset, we randomly sample 50,000 sentences, partitioning them into 40,000 for training, 5,000 for validation, and 5,000 for testing.

5.2 Models

We employ two pre-trained language models: BERT-base¹ and BERT-large² (cased) (Devlin et al., 2019). BERT-base uses 12 layers, 12 heads, and a 768-dimensional hidden state, while BERT-large uses 24 layers, 16 heads, and a 1024-dimensional hidden state. These models provide a range of capacities, allowing us to investigate differences in how syntactic structures are constructed across models.

We specifically focus on BERT models because our method is designed to examine the atemporal, layer-wise derivation of syntactic structures after observing an entire sentence. In contrast, autoregressive models like GPT-2 inherently construct syntactic information incrementally in a left-to-right manner, a fundamentally different setting from the layer-wise derivation we aim to study. Nevertheless, our method is also applicable to GPT-2, and we report GPT-2 results separately in App. B.

For each model, we probe all layers to determine the progression of syntactic information and compute the expected layer at which specific structures emerge. We conduct training with five different random seeds and report the average performance

¹<https://huggingface.co/google-bert/bert-base-cased>

²<https://huggingface.co/google-bert/bert-large-cased>




Figure 3: Global UUAS by each layer for each model. Error bars represent standard deviations across 5 random seeds.

along with the standard deviation. Additional hyperparameters and training details are provided in App. C.

6 Results

6.1 Overall UUAS Performance

As a sanity check to verify whether our models exhibit overall trends similar to those reported in previous studies, we conducted an experiment measuring the test set UUAS for overall sentence structures across layers for each model (Figure 3). BERT-base and BERT-large display similar trends, with the UUAS score saturating around the middle layers. BERT-large shows slightly slower improvement, likely reflecting its deeper architecture and larger capacity. These trends mostly align with previous findings (Hewitt and Manning, 2019) that neural language models tend to exhibit peak UUAS performance in their middle layers. However, unlike previous studies, we do not observe a decrease in average UUAS in later layers, which we attribute to our method of computing word embeddings as a weighted average from layer 0 to layer k (Eq. (3)).

6.2 Expected Layer Across Structure Sets

Figure 4 summarizes the expected layers for each syntactic structure within the four primary structure sets (§5.1), for both BERT-base and BERT-large.

In both BERT-base and BERT-large, the macro-syntactic structure consistently exhibits the highest expected layer across all sets, whereas micro-syntactic structures such as *nsubj*, *dobj*, and *prep* tend to appear in lower layers. This suggests a *bottom-up* derivation process in which




Figure 4: Expected layer for each model across different structure sets. Error bars represent standard deviation across 5 random seeds.

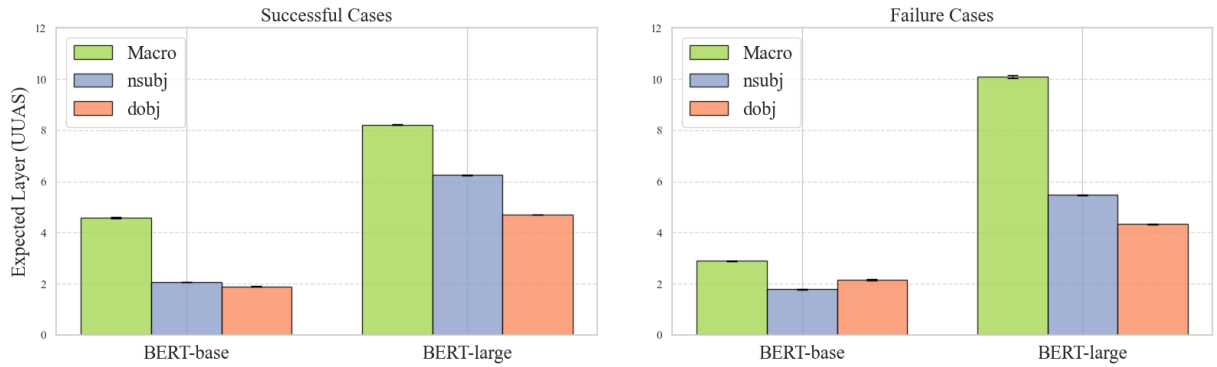


Figure 5: Expected layers for syntactic structures in successful and failed subject-verb agreement cases. Error bars show standard deviations across 5 random seeds.

micro-syntactic structures (e.g., subject or object phrases) are constructed earlier, and these components are gradually integrated into a coherent macro-syntactic structure in later layers. This observation is consistent with prior work on BERT showing that local information (e.g., POS tags) is captured early while more abstract global structures emerge later (cf. Tenney et al., 2019). Our results extend these findings by demonstrating that BERT builds the full syntactic structure by first constructing micro-syntactic structures before assembling the macro-syntactic structure. Notably, this pattern holds for both BERT-base and BERT-large, although the overall expected layers are slightly higher in BERT-large—likely reflecting its deeper architecture and larger capacity.

7 Detailed Analysis: Subject-Verb Agreement Task

7.1 Experimental Setup

To investigate how the process and layers involved in syntactic structure construction relate to model performance, we conduct a detailed analysis on subject-verb agreement using sentences with intervening nouns (“attractors”), following the approach of Marvin and Linzen (2018) with some modification. We sampled 1,000 positive (grammatical) and 1,000 negative (ungrammatical) sentences. In our modification of their templates, each of the subject noun phrase, prepositional phrase, and verbal phrase is required to contain more than one words. This ensures that we can extract meaning-

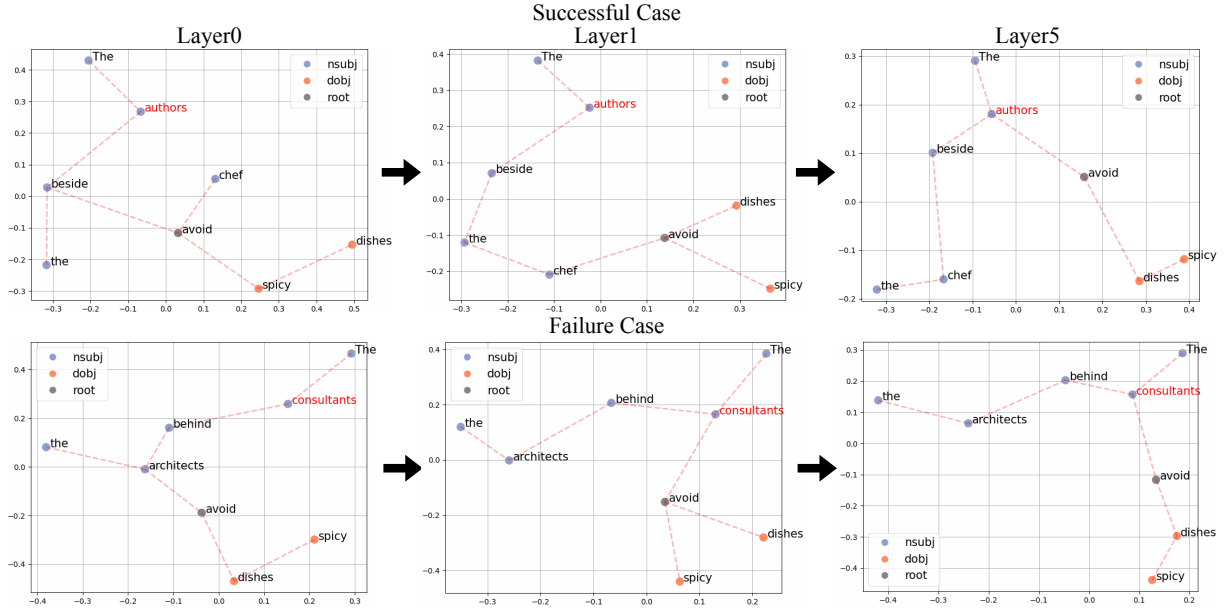


Figure 6: Derivation process visualizations for BERT-base on subject-verb agreement for a successful case (“The surgeon behind the guards attends exclusive art exhibitions.”) and an failure case (“The consultants behind the architects avoid spicy dishes.”).

ful subgraphs within each syntactic substructure. For instance, our dataset includes the following sentences:

- a. The senators behind the brilliant architect avoid spicy dishes.
- b. *The senators behind the brilliant architect avoids spicy dishes.

We first evaluate model performance on this task by computing pseudo-whole-sentence probabilities. Specifically, we calculate the probability of each token by masking it one by one and then aggregate these token-level probabilities to derive an overall sentence probability. We expect the model to assign higher pseudo-probabilities to grammatical sentences compared to ungrammatical ones. We then analyze how the syntactic construction process differs between cases where the model performs well and those where it fails.

Furthermore, to qualitatively visualize the evolution of syntactic structures across model layers, we employ Multidimensional Scaling (MDS). Specifically, we apply scikit-learn’s MDS implementation (Pedregosa et al., 2011) with default parameters to word embeddings projected by our structural probe, allowing us to illustrate clearly how syntactic representations develop across different model layers.

7.2 Results

Overall, BERT-base correctly assigned higher pseudo-whole-sentence probabilities to grammatical sentences in 984 out of 1,000 examples, whereas BERT-large achieved correctness in 983 out of 1,000 cases. Despite their similar overall accuracies, we observe distinct patterns between BERT-base and BERT-large (Figure 5).

BERT-base. BERT-base frequently failed when macro-syntactic structures were established prematurely, potentially restricting the incorporation of essential micro-syntactic details. As illustrated in Figure 6, successful cases show a sequential pattern where BERT-base first constructs micro-syntactic structures within the subject phrase in early layers, subsequently aligning the subject (*authors*) with the verb (*avoid*) around layer 5 after stabilizing the internal subject dependencies. In contrast, failure cases reveal premature alignment of macro-syntactic structures, with the subject (*consultants*) prematurely linked to the verb (*avoid*) before fully establishing necessary micro-syntactic details. This premature commitment might have negatively impacted the overall syntactic representation, disrupting correct subject-verb agreement.

BERT-large. BERT-large exhibited higher expected layers for macro-syntactic structures in failure cases, suggesting delayed integration of macro-

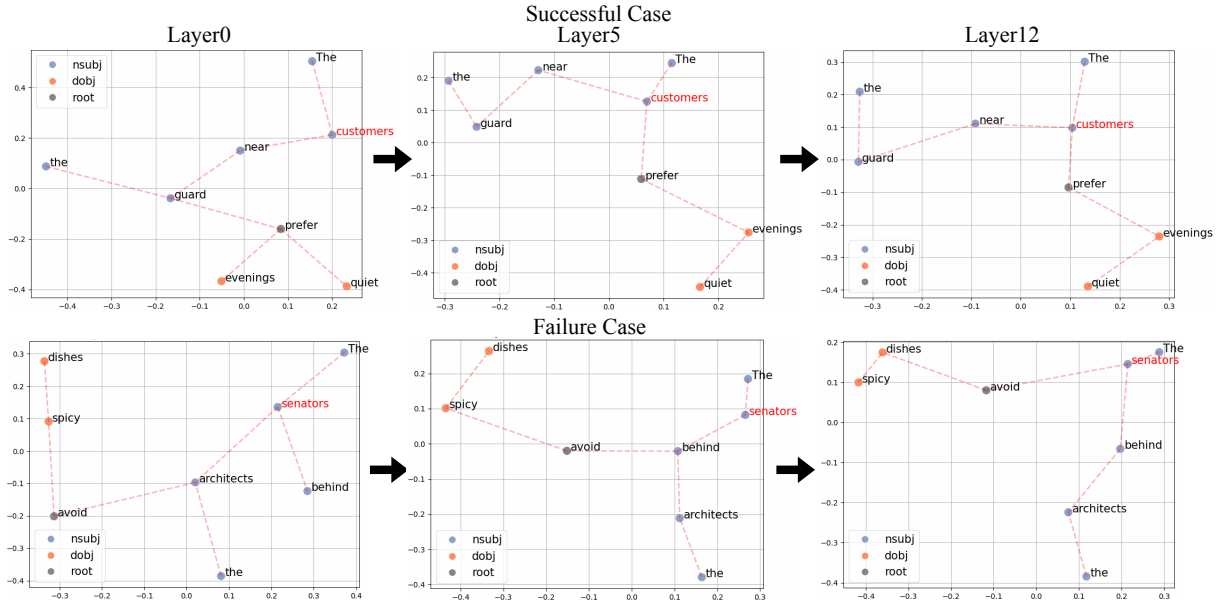


Figure 7: MDS visualizations of syntactic structure evolution in BERT-large for subject-verb agreement for a successful case (“The customers near the guard prefer quiet evenings.”) and an failure case (“The senators behind the architects avoid spicy dishes.”). Red highlights indicate the correct subject.

syntactic information. Figure 7 illustrates representative successful and unsuccessful cases for BERT-large. Successful predictions demonstrate early alignment of the subject (*customers*; highlighted in red) with the verb (*prefer*) around layer 5, facilitating accurate subject-verb agreement. Conversely, in unsuccessful cases, this alignment emerged considerably later (around layer 12), highlighting delayed macro-syntactic integration.

Taken together, these analyses suggest an optimal intermediate range of layers for integrating macro-syntactic information. Forming macro-syntactic structures either prematurely or excessively late can negatively affect syntactic processing, highlighting the importance of appropriately timed integration for accurate predictions. These visualizations underscore how deviations from this optimal timing contribute to subject-verb agreement errors.

8 Discussion and Conclusion

In this paper, we introduced *Derivational Probing*—a method that integrates structural probing with an expected layer metric to trace the construction process of syntactic structures in neural language models. Our experiments revealed that BERT models tend to build micro-syntactic dependencies first and gradually assemble them into a coherent macro-structure.

BERT’s bidirectional context supports a step-

wise, bottom-up construction—starting with the formation of local, micro-syntactic structures and culminating in a fully integrated macro representation. These findings offer valuable insights into the internal mechanisms by which deep neural models construct syntactic trees and highlight the importance of examining layer-wise structural formation for improved model interpretability.

One promising avenue for future research is the design of more controlled probing tasks and datasets. By constructing minimal pairs that differ only in specific grammatical properties (or by using Jabberwocky-style sentences devoid of semantics; Maudslay and Cotterell, 2021), future studies can better isolate the abstract syntactic rules that models learn. Moreover, incorporating multilingual and cross-domain probes will help determine whether these syntactic features generalize beyond English or are artifacts of a particular corpus. Finally, another exciting direction is to explore left-to-right parsing strategies as an alternative to tracking the construction process through layers (e.g., Eisape et al., 2022), which could yield further insights into the role of processing order in syntactic representation.

References

Tyler A. Chang and Benjamin K. Bergen. 2024. *Language Model Behavior: A Comprehensive Survey*. *Computational Linguistics*, 50(1):293–350.

Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What does BERT look at? an analysis of BERT's attention . In <i>Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP</i> , pages 276–286, Florence, Italy. Association for Computational Linguistics.	584
Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding . In <i>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)</i> , pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.	585
Tiwalayo Eisape, Vineet Gangireddy, Roger Levy, and Yoon Kim. 2022. Probing for incremental parse states in autoregressive language models . In <i>Findings of the Association for Computational Linguistics: EMNLP 2022</i> , pages 2801–2813, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.	586
Katsuhiko Hayashi, Taro Watanabe, Masayuki Asahara, and Yuji Matsumoto. 2012. Head-driven transition-based parsing with top-down prediction . In <i>Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 657–665, Jeju Island, Korea. Association for Computational Linguistics.	587
John Hewitt, Kawin Ethayarajh, Percy Liang, and Christopher Manning. 2021. Conditional probing: measuring usable information beyond a baseline . In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 1626–1639, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	588
John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations . In <i>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)</i> , pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.	589
Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python .	590
Jenny Kunz and Marco Kuhlmann. 2022. Where does linguistic information emerge in neural language models? measuring gains and contributions across layers . In <i>Proceedings of the 29th International Conference on Computational Linguistics</i> , pages 4664–4676, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.	591
Tomasz Limisiewicz and David Mareček. 2021. Introducing orthogonal constraint in structural probes . In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 428–442, Online. Association for Computational Linguistics.	592
Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. In <i>Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing</i> , pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.	593
Rowan Hall Maudslay and Ryan Cotterell. 2021. Do syntactic probes probe syntax? experiments with jabberwocky probing . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 124–131, Online. Association for Computational Linguistics.	594
Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models . <i>Preprint</i> , arXiv:1609.07843.	595
Joakim Nivre. 2004. Incrementality in deterministic dependency parsing . In <i>Proceedings of the Workshop on Incremental Parsing: Bringing Engineering and Cognition Together</i> , pages 50–57, Barcelona, Spain. Association for Computational Linguistics.	596
F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. <i>Journal of Machine Learning Research</i> , 12:2825–2830.	597
Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations . In <i>Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)</i> , pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.	598
R. C. Prim. 1957. Shortest connection networks and some generalizations . <i>Bell System Technical Journal</i> , 36(6):1389–1401.	599
Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT rediscovers the classical NLP pipeline . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 4593–4601, Florence, Italy. Association for Computational Linguistics.	600
Jesse Vig and Yonatan Belinkov. 2019. Analyzing the structure of attention in a transformer language model . In <i>Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP</i> , pages 63–76, Florence, Italy. Association for Computational Linguistics.	601

- Jennifer C. White, Tiago Pimentel, Naomi Saphra, and Ryan Cotterell. 2021. [A non-linear structural probe](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 132–138, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. [Explainability for large language models: A survey](#). *ACM Trans. Intell. Syst. Technol.*, 15(2).