

What if This Modified That? Syntactic Interventions via Counterfactual Embeddings

Mycal Tucker
MIT
mycal@mit.edu

Peng Qian
MIT
pqian@mit.edu

Roger P. Levy
MIT
rplevy@mit.edu

Abstract

Neural language models exhibit impressive performance on a variety of tasks, but their internal reasoning may be difficult to understand. Prior art aims to uncover meaningful properties within model representations via probes, but it is unclear how faithfully such probes portray information that the models actually use. To overcome such limitations, we propose a technique, inspired by causal analysis, for generating counterfactual embeddings within models. In experiments testing our technique, we produce evidence that suggests some BERT-based models use a tree-distance-like representation of syntax in downstream prediction tasks.

1 Introduction

Large neural models like BERT and GPT-3 have established a new state of the art in a variety of challenging linguistic tasks (Devlin et al., 2019; Brown et al., 2020). These connectionist models, trained on large corpora in a largely unsupervised manner, learn to map words into numerical representations, or embeddings, that support language-reasoning tasks. Fine-tuning these models on tasks like extractive question answering specializes these generic models into performant, task-specific models (Wolf et al., 2019).

In conjunction with the rise of these powerful neural models, researchers have investigated what the models have learned. Probes, tools built to reveal properties of a trained model, are a favored approach (Hall Maudslay et al., 2020; Conneau et al., 2018). For example, Hewitt and Manning (2019) have uncovered compelling evidence that several models encode syntactic information in their embeddings. That is, by passing embeddings through a trained probe, one may recover information about a sentence’s syntax.

“I saw the boy and the girl [MASK] tall.”

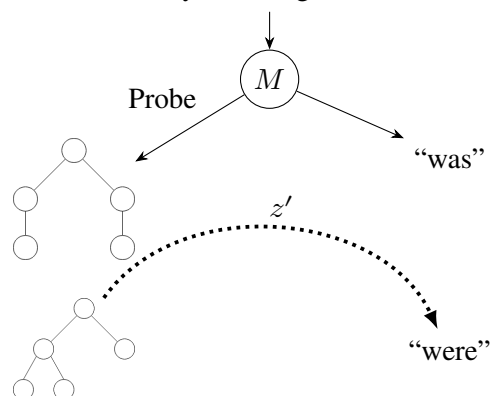


Figure 1: A language model, M , outputs predictions and a probe estimates properties from the model representation. We use probes to generate counterfactual representations, z' , based on syntactic manipulations, revealing reasoning within the model.

Although these results are impressive, they fall short of clearly demonstrating what linguistic information the language models actually use. Syntactic information is present in sentences; that embeddings also encode syntax does not imply that a model uses syntactic knowledge.

In order to truly query a model’s understanding, one must use causal analysis. Recently, several authors have done so by generating counterfactual data to test models (Kaushik et al., 2020; Goyal et al., 2019; Elazar et al., 2020). They either create new input data or ablate parts of embeddings and study how model outputs change. We extend this prior art via a new technique for generating counterfactual embeddings by using traditional probes to manipulate embeddings according to syntactic principles, as depicted in Figure 1. Because we conduct experiments with syntactically ambiguous inputs, we are able to measure how models respond to different valid parses of the same sentence instead of, for example, removing all syntactic information.

Thus, our technique uncovers not only what parts of its embeddings a model uses to represent syntax, but also how those parts influence downstream behavior.

Thus, in this work, we make two contributions. First, we develop a gradient-based algorithm to generate counterfactual embeddings, informed by trained probes. Second, in experiments using our technique, we find that the standard BERT model, trained on word-masking tasks, appears to leverage features of syntax in predicting masked words but that a BERT model fine-tuned for question-answering does not. In addition, these experiments yield new data to inform the ongoing debate on probe design.

2 Related Work

2.1 Neural Language Model Probes

Transformer-based models like GPT-3 and BERT have recently advanced the state of the art in numerous language-related problems (Brown et al., 2020; Devlin et al., 2019; Wolf et al., 2019). These large models appear to learn meaningful representations of words and sentences, enabling high performance when fine-tuned for a specific task.

In conjunction with these models, probes have been developed to uncover what principles models have learned. Such probes have been used in a wide variety of contexts, from image structure to syntax and semantics in language models (Alain and Bengio, 2018; Conneau et al., 2018; Hewitt and Manning, 2019; Coenen et al., 2019, among others). Our work uses two syntactic probes developed by Hewitt and Manning (2019) that map from model embeddings to predictions about word locations in a parse tree. These probes are simple by design – merely linear transformations – in order to prevent the probes themselves from doing parsing.

Recent work directly addresses the topic of probe simplicity. On the one hand, if probes are too expressive, they may reveal their own learning instead of a model’s (Liu et al., 2019; Hewitt and Liang, 2019). On the other hand, Pimentel et al. (2020) argue from an information-theoretical perspective that more expressive probes are always preferable.

Our work differs from much prior art in probe design by leveraging causal analysis, which uses counterfactual data to test probes and models. This provides direct evidence of whether a model uses the same features as a probe, allowing us to experi-

ment beyond linear probes (and indeed, we found that more complex probes offered an advantage in some cases).

2.2 Causal Analysis of Language Models

Motivated by the limitations of traditional, correlative probes, researchers have recently turned to causal analysis to better understand language models. Goyal et al. (2019) and Kaushik et al. (2020) generate counterfactual inputs to language models, while Vig et al. (2020) study individual neurons and attention heads to uncover gender biases in pre-trained networks.

Our work is most closely related to that of Elazar et al. (2020), who, as in this work, used probes to generate counterfactual embeddings within a network. Their amnesiac counterfactuals are generated by suppressing features in embeddings that a probe uses. In contrast, we use a continuous, gradient-based approach to generate counterfactuals, yielding insight into how features are used, as opposed to if they are used at all.

3 Technical Approach

3.1 Problem Formulation

We may characterize a transformer-based language model, M , trained on a specific task, as a function mapping from an input string, s , to an output y : $M(s) = y$. In order to reveal embeddings for analysis by probes, we may decompose M into two functions: M_{k-} and M_{k+} . M_{k-} represents the first k layers of the model; M_{k+} represents the layers of M after layer k ; M is the composition of these functions: $M = M_{k+} \circ M_{k-}$. We label the embeddings output by M_{k-} as z_k . This decomposition of models to reveal internal embeddings mirrors the formulation for layer-specific probes (Hewitt and Manning, 2019). A probe may be defined as a function f_p that maps from an embedding, z_k , to a predicted property \hat{p} about the input, s : $f_p(M_{k-}(s)) = \hat{p}$. (For the remainder of this paper, we focus on syntactic probes, but our reasoning may be extended to other linguistic properties.)

We may define two, potentially overlapping, subsets of the features of z_k by considering different uses of z_k . First, we may define z_p as the features of z_k that the probe uses in predicting \hat{p} (for example, when using a linear probe, z_p is the projection of z_k onto the probe subspace). Assuming good syntactic probe performance, z_p is necessarily informative of the input’s syntax. We likewise

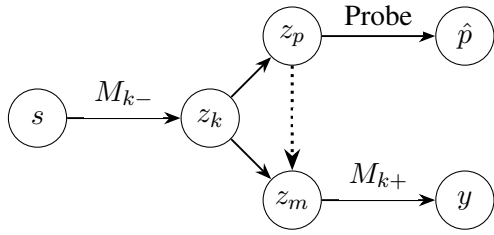


Figure 2: M_{k-} yields a representation, z_k . z_p and z_m are subsets of the features of z_k used by the probe and M_{k+} . We measured the causal link between z_p and z_m .

defined z_m as the features of z_k that M_{k+} uses in producing the model output. These two, potentially overlapping, representations of z_k are shown in Figure 2, inspired by causal diagrams by Pearl and Mackenzie (2018). We seek to discover if there is a causal link between z_p and z_m .

For some tasks, such a link should exist. For example, a question-answering model’s response to “I shot the elephant wearing my pajamas. Who wore the pajamas?” should depend upon the inferred sentence syntax (e.g., if the probe predicts that “wearing my pajamas” modifies “the elephant,” the model should output “the elephant”). Thus, the probe and model outputs should “agree” according to syntactic principles. Furthermore, if a causal link between z_p and z_m exists, changing z to produce a new prediction of syntax should change the model output to agree with the probe (e.g., if the probe predicts that “wearing my pajamas” now modifies “I,” the model should now output “I”). In this work, therefore, we study whether a link between z_p and z_m exists and, if it does, to what extent it corresponds with linguistic principles.

3.2 Generating Counterfactual Embeddings via Gradient Descent

To study such a link, we must generate counterfactual embeddings, z' , that modify probe outputs, starting from normal embeddings z_k . We borrow the term “counterfactual” from causal literature because z' represents what z_k would have been if z_p had been different (Pearl and Mackenzie, 2018). We were particularly interested in finding z' that changed both probe and model outputs; if z' only changed probe outputs, that could indicate that the probe was over-interpreting model embeddings (e.g., acting as a parser instead of a probe).¹

We developed a gradient-based method to gener-

¹We did not study z' that only modified the model outputs, although this could be a promising avenue for future work.

ate z' that changed the probe output. We assumed that, given the probe function, f_p , a loss, L , and the correct property value (e.g., parse), p , one could compute the gradient of the loss with respect to the probe inputs: $\nabla_{z'} L(f_p(z'), p)$. Neural network probes obey such differentiability assumptions.

Given z_k and p , we constructed a counterfactual embedding, z' , by initializing z' as a z_k generated by the model and updating z' via gradient descent of the loss. Updating z' may be terminated based on various stopping criteria (e.g., local optimality, loss below a threshold, etc.), yielding the final counterfactual z' . Assuming non-zero gradients, this technique produces z' s that, by design, change the probe outputs. In experiments, we studied how z' s changed model outputs when passed through M_{k+} .

Although our technique bears some resemblance to gradient-based adversarial attacks (Szegedy et al., 2014), it may more broadly be thought of as guided search in a latent space. Adversarial images are often characterized by changes that are imperceptible to humans but change model behaviors to be incorrect. In contrast, we seek to find embeddings that change both probe and language model outputs. Furthermore, by design, we use syntactically ambiguous sentences in experiments and generate counterfactuals according to valid parses. Thus, unlike adversarial attacks on images that seek to switch model classification to an incorrect class, we merely guide embeddings among a set of valid interpretations. Lastly, even uncovering instances of embeddings that change probe outputs but not the model’s is important, as it indicates a misalignment of probe and model reasoning.

4 Experiments

In the previous section, we proposed a technique for generating counterfactual embeddings; here, we detailed the experiments we conducted to measure the effect of using such embeddings. Inputs to our technique included the base language models, probes, test sentences, and different ground-truth parses to generate the counterfactual embeddings.

4.1 Model Tasks

We tested our technique on two BERT models trained on different tasks: masked word prediction and extractive question answering.

In the masked word prediction task, a model is given a sentence, S , comprising words $(s_0, s_1, \dots$

Model	Corpus	Parse	Example Input
Mask	Coord.	Plur. Sing.	The woman saw (the boy and the dog) [MASK] falling. (The woman saw the boy) and (the dog [MASK] falling.)
	NP/Z	Adv. Noun	When the dog scratched (the vet [MASK] ran.) (When the dog scratched the vet) [MASK] ran.
QA	RC	Conj. NP2	The ((smart women and rich men) who were desperate) bribed the judge. The (smart women) and (rich men who were desperate) bribed the judge.
	NP/VP	VP NP2	The girl saw (the boy) with the telescope. The girl saw (the boy with the telescope.)

Table 1: Experiment design for different language models and test corpora, with illustrative sentences, decorated with auxiliary parentheses to reveal structure. The parentheses were not included in the actual corpora.

[MASK], ..., s_n) and must predict the word at the location marked by [MASK]. For example, given a sentence, [‘The’, ‘children’, ‘went’, ‘out’, ‘to’, [MASK], ‘.’], a correct answer might be ‘play.’ We used huggingface’s “bert-large-uncased-whole-word-masking” model, which was trained on masked word and next-sentence prediction, and referred to it as the “Mask” model (Wolf et al., 2019).

Extractive question answering is framed by Wolf et al. (2019) as follows: given a sentence, S , comprising word tokens (s_0, s_1, \dots, s_n) and a question, identify the start and end tokens ($s_i, s_j; 0 \leq i \leq j \leq n$) denoting a contiguous stretch of the sentence that answers the question. For example, given the sentence [‘I’, ‘ate’, ‘two’, ‘apples’, ‘.’] and the question “How many apples did I eat?,” a correct answer could be [2, 2] (“two”) or [2, 3] (“two apples”). We used huggingface’s “BertForQuestionAnswering,” already fine-tuned on the SQuAD dataset, and referred to the model as QA (Wolf et al., 2019; Rajpurkar et al., 2016).

4.2 Probes

Our technique for generating counterfactual embeddings depended on probes, so we used four different syntactic probes drawn from prior art and our own design.

The depth probe from Hewitt and Manning (2019) maps from embeddings to predictions over words’ depths in a sentence’s parse tree. The distance probe, given a pair of words, predicts the distance between the words in the parse tree (i.e., how many edges must be traversed). Both probes consist of a linear transformation from embedding to prediction.

We further implemented “deep” versions of the distance probe by creating two- and three-layer,

non-linear probes trained on the distance task. These models used ReLU activations, with hidden dimension 1024, but otherwise used the same input and output format as the linear distance probe. (Experiments conducted with “deep” versions of the linear depth probe produced similar results to those of the normal depth probe and are therefore omitted.)

4.3 Evaluation Corpora

We used four corpora for evaluating the Mask and QA models, as summarized in Table 1.

4.3.1 Mask Test Corpora

For the Mask model, we used two test suites composed of sentences whose structural ambiguity was resolved by a masked word.

The first corpus, dubbed “Coordination,” comprised sentences that took the form “The NN1 VERB the NN2 and the NN3 [MASK] ADJ.” Such sentences may be interpreted in at least two ways by inserting either “was” or “were” in the masked location. The former reflects a conjunction of clauses (e.g., “The woman saw the boy and the dog was falling.”), whereas the latter reflects a conjunction of noun phrases (e.g., “The woman saw the boy and the dog were falling.”) Sentences were generated through combinations of NN1 [man, woman, child], VERB [saw, feared, heard], NN2 [boy, building, cat], NN3 [dog, girl, truck], and ADJ [tall, falling, orange], yielding 243 sentences, each with two parse trees dubbed “singular” or “plural,” depending on the grammatical verb type.

The second corpus, dubbed the NP/Z corpus, was inspired by classic psycholinguistic studies of the garden-pathing effect in online sentence processing (Frazier and Rayner, 1982; Tabor and Hutchins, 2004). Each sentence in the corpus took the form “When the NN1 VERB1 the NN2

[MASK] VERB2.” Without knowing the masked word, it is unclear if NN2 is the object of the subordinate clause or the subject of the main clause. For example, in the sentence “When the dog scratched the vet [MASK] ran,” either an adverb (e.g., “immediately”) or a noun (e.g., “she”) would be permitted but correspond to different parses. We created such parse trees and dubbed the first type “Adv.” and the second type “Noun.” We used the 24 sentences from [Tabor and Hutchins \(2004\)](#) that fit our template, and supplemented the dataset with 36 sentences of our own, generated by iterating over all combinations of NN1 [dog, child], NN2 [vet, boy, girl], VERB1 [scratched, bit], and VERB2 [ran, screamed, smiled]. (Augmenting the dataset was needed to increase the statistical analysis power, and plotting the 24 and 36 sentences separately established that they produced similar results.)

4.3.2 QA Test Corpora

For the QA model, we created two test suites. First, the “RC” corpus used sentences composed of a conjunction of nouns modified by a relative clause. All sentences took the form “The ADJ1 NN1 and ADJ2 NN2 who were ADJ3 VERB the NN3. Who was ADJ3?” For example, one sentence was “The smart women and rich men who were desperate bribed the judge. Who was desperate?” By construction, it was unclear if the relative clause modified the conjunction of the first and second noun phrases (The ADJ1 NN1 and ADJ2 NN2) or merely the second noun phrase (ADJ2 NN2). For each sentence, we generated two parses: “Conj. Parse” and “NP2 Parse,” corresponding to the former and latter. We generated sentences by iterating over all combinations of values for ADJ1 [smart, rich, tall, poor], NN1 [men, women], ADJ2 [smart, rich, tall, poor], NN2 [men, women], ADJ3 [corrupt, desperate], VERB [bribed, paid], and NN3 [politician, judge], excluding sentences in which NN1 and NN2 or ADJ1 and ADJ2 were the same. This produced 192 sentences, each with two parses.

Lastly, the “NP/VP” corpus used sentences with ambiguous prepositional phrase attachment. Inspired by sentences like “The girl saw the boy with the telescope,” we generated inputs with the template “The NN1 VERB the NN2 with the NN3. Who had the NN3?” We iterated through combinations of NN1 [man, woman, child], NN2 [man, woman, boy, girl, stranger, dog], and VERB-NN3 pairs [saw-telescope, poked-stick, thanked-letter, fought-knife, dressed-hat, indicated-ruler, kicked-

shoe, welcomed-gift, buried-shovel], removing duplicate NN1 and NN2, yielding 144 inputs. Each input used two parses indicating the prepositional phrase modifying VP or NP2 (“the” and NN2).

4.4 Generating Embeddings

For all models, probes, and parse trees for each sentence, we generated counterfactual embeddings by initializing a counterfactual embedding, z' , as the original model embedding for the input sentence, z_k , and running an Adam optimizer, with learning rate 0.0001, to minimize the probe loss (using a particular probe and parse tree) ([Kingma and Ba, 2014](#)). Recall that the optimizer updated z' rather than the probe parameters.

The optimizer used a patience value of 5000: it continued updating z' until the probe loss failed to improve for 5000 consecutive gradient updates. Using a patience-based termination condition (as opposed to setting a loss threshold or maximum number of updates, for example) was task-agnostic and seemed to be robust to a wide range of patience values. Brief experimentation with patience values from 50 to 5000 yielded similar results. On a Linux desktop with an Nvidia GeForce RTX 2080 graphics card, generating a single counterfactual took less than 1 minute, and the process was easily parallelized to batches of 80 embeddings, reducing the mean computation time to under one second.

For both the QA and Mask models, we trained all probe types (depth, distance, 2-layer dist, and 3-layer dist) on each of the model’s 25 layers. We used 5000 entries from the Penn Treebank (PTB) for training, with the standard validation and test sets of nearly 4000 entries used for early stopping and evaluation, respectively ([Marcus et al., 1993](#)).

4.5 Metrics

We used two sets of metrics in our experiments. First, we measured probe performance using the Root Accuracy, UUAS, and Spearman Coefficient metrics used by [Hewitt and Manning \(2019\)](#) and refer to their work for details. Intuitively, these metrics captured how accurately the probes predicted aspects of syntactic structure from embeddings.

Second, we measured changes in model outputs when using counterfactual embeddings. The Mask model produced a probability distribution over more than 30,000 possible words for the masked location, but we restricted our attention to only a subset of those words, dubbed “candidates.” (We normalized predictions among the set of candidates,

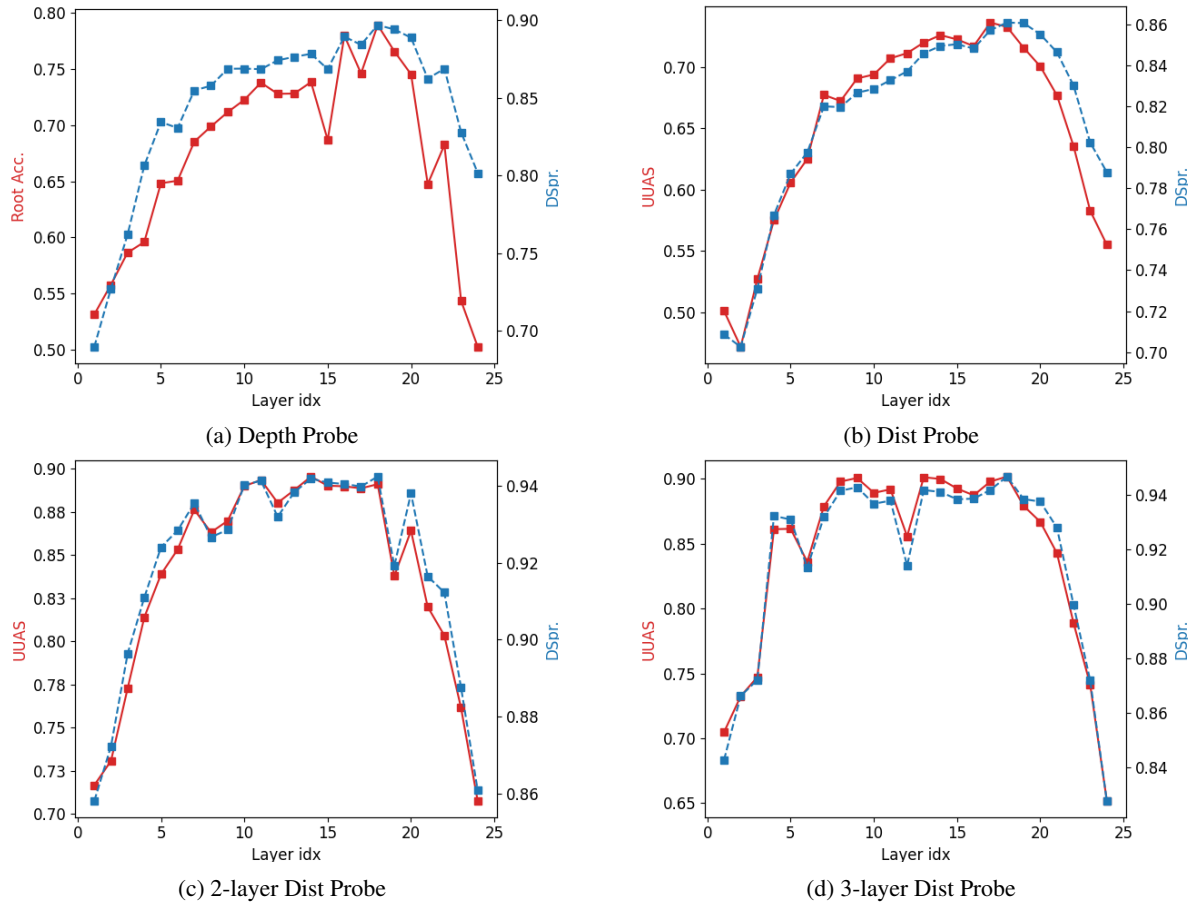


Figure 3: All trained probes for the QA model exhibited high performance on the PTB corpus.

producing a proper probability distribution.) In the Coordination corpus, we used 5 candidates: [“was,” “is,” “were,” “are,” “as”]. In the NP/Z corpus, we generated the set of candidates by collecting the most likely predictions over the corpus, using both original and counterfactual embeddings. This set of 18 words is shown in the x -axis of Figure 6. For both corpora, we partitioned the candidates into two sets, depending upon which parse they implied, and measured the sum of the probabilities of words in each set. If counterfactual embeddings caused the models to change the type of word they predicted, we would expect to see a change in these sums.

For the QA model, we similarly measured changes in probabilities among sets of words, but in this case we focused on the predicted start location of the answer. Recall that the QA model produced two distributions over words, indicating its predictions over where the answer started and ended. Consider an example input, drawn from the RC corpus: “The smart women and rich men who were desperate bribed the politician. Who

was desperate?” Two reasonable answers might be “The smart women and rich men” or “rich men,” corresponding to QA outputs with identical end words, but differing start words. We therefore created two partitions of starting words to consider: those belonging to the first noun phrase (“The smart women”) or the second noun phrase (“rich men”). We then measured the summed start probabilities of words in each partition. We did not normalize these probabilities, as the QA model rarely predicted start words outside these two partitions with more than 1% probability.

In all experiments, we employed one-sided Wilcox Signed-Rank tests, non-parametric tests for pairwise data, when determining the significance at ($p < 0.01$). The parses were viewed as “treatments” for the same embedding. We compared the effect of using counterfactual instead of original embeddings, as well as the effect of using different parses to generate counterfactual embeddings.

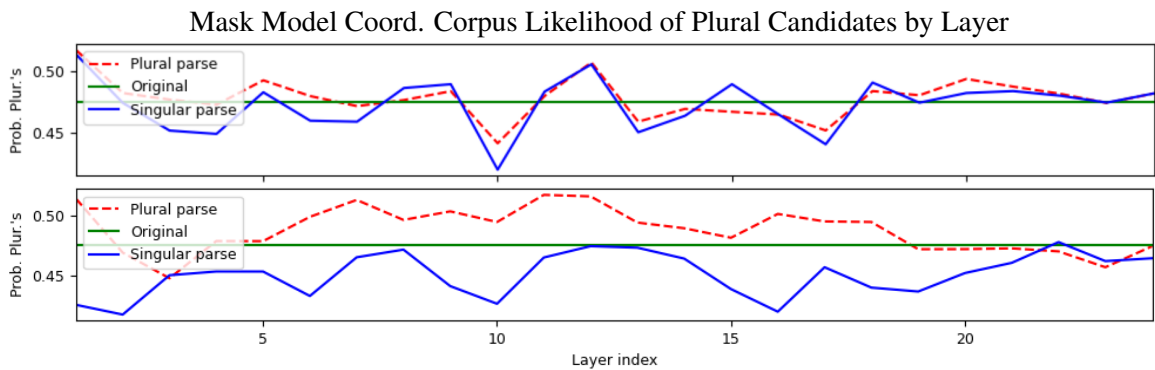


Figure 4: Mean probability of plural candidates using the depth probe (top) or the 3-layer dist probe (bottom), using original or counterfactual embeddings, in the Coordination corpus. Using a parse that implied plural words increased the probability of plural words when using the 3-layer dist probe.

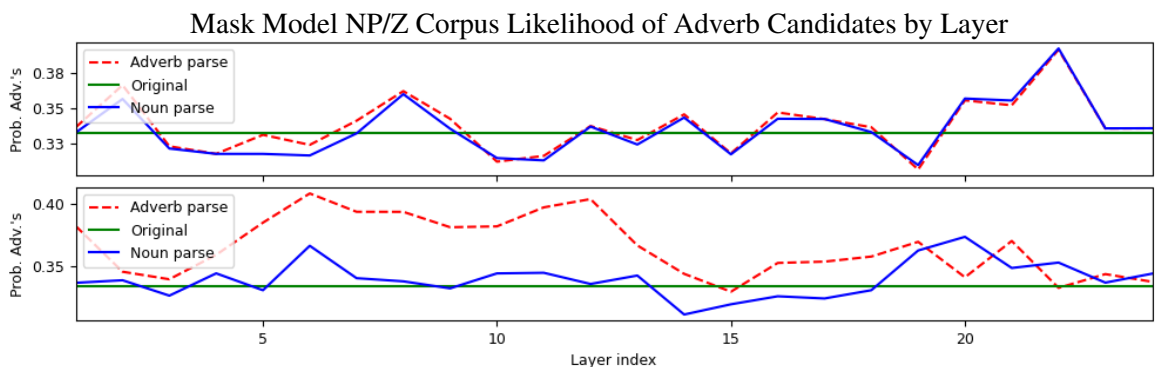


Figure 5: Mean probability of adverb candidates in the NP/Z corpus, using original and counterfactual embeddings generated by the depth (top) and 3-layer dist probes (bottom).

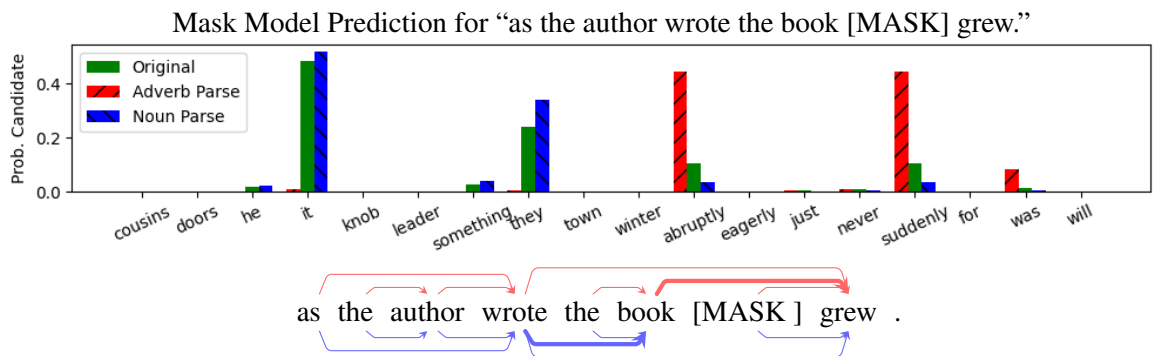


Figure 6: Given a sentence from the NP/Z corpus, the Mask model originally predicted “it” or “they,” but using counterfactuals from the 5th layer 3-layer dist probe changed predictions to favor nouns (cousins - winter) or adverbs (abruptly - suddenly). Visualizing the word dependencies revealed that the Adverb parse (top, red) and Noun parse (bottom, blue) induced different dependencies (differences in bold), as expected.

5 Results

Our results indicated that our probes performed well, as evaluated by performance metrics from prior art. However, we found that only some combinations of probe types and BERT models generated counterfactuals that altered the model’s outputs according to syntactic principles.

5.1 Probe Performances

Measured on the PTB test set, the probe performance metrics confirmed that the probes predicted aspects of syntactic structure well (Marcus et al., 1993). Plots of performance, similar to those by Hewitt and Manning (2019), for probes trained on QA model embeddings are included in Figure 3.²

²All probe metrics are plotted in the appendix.

For both models and all probe types, we found that the probes were able to achieve high performance, indicating that both the Mask and QA models encoded syntactic information in their embeddings.

We also observed the unsurprising trend that multi-layered, non-linear distance probes outperformed the linear distance probe. This raised the question, if different probes exhibited different performance for the same model, which probe should be used to deduce model behavior? Injecting counterfactual embeddings generated by different probes helped us answer this question.

5.2 Mask Counterfactual Results

Next, we found that using the distance-based probes to generate counterfactual embeddings in the Mask model consistently produced the desired effect by shifting the model’s prediction of the masked word according to syntactic principles, and that the multi-layer distance probes performed better than the linear probe.

We plotted the mean effect of counterfactual embeddings for the Coordination and NP/Z corpora in Figures 4 and 5, respectively.³ Each plot depicts the mean prediction likelihood of one of the partitions of candidates (plural for Coord. corpus, adverbs for NP/Z), using original or counterfactual embeddings. Figure 4 shows results using the depth and 3-layer distance probes in the Coord. corpus: the depth probe failed to produce consistent changes in word probabilities, but embeddings generated by the 3-layer dist probe did exhibit the desired effect. The change in probability of plural words when using the plural parse was significantly positive for layers 6 through 14 (among others) and greater than the change when using the singular parse for layers 4 through 21.

Similar results were observed using the 3-layer distance probe for the NP/Z corpus, as shown in Figure 5. The net increase in probability for adverbs when using the adverb parse was significantly greater than when using the NP2 parse for layers 5 through 19 and was positive for layers 4 through 13.

We examined an example sentence from the NP/Z corpus in Figure 6 in greater depth. The 18 words displayed along the x axis were the candidate words whose probabilities we calculated in the NP/Z corpus. As expected, using the Adv. parse

increased the likelihood of adverbs like “suddenly,” while using the Noun parse increased the likelihood of nouns like “it” or “they.” Lastly, the bottom part of Figure 6 shows the dependency trees for the counterfactuals generated for each parse (see Hewitt and Manning (2019) for details on creating such trees). These trees reflected the dependencies of the parses that generated the counterfactuals, indicating that our technique changed embeddings in the way we intended.

Together, the results from both corpora, revealed that distance-, but not depth-, based probes elicited the desired response from the Mask model, which suggests that it leverages a distance-based representation of syntax in its reasoning.

5.3 QA Counterfactual Results

Lastly, we examined the effect of using counterfactual embeddings in the QA model. Compared to the Mask model, we found smaller and less consistent results, suggesting that the QA model may not use syntax.

Taking the mean across sentences in the corpus, we plotted the mean starting probabilities of words in each sentence’s first noun phrase (as explained earlier in Section 4.5). These values reflect whether the model predicted NP1 should be included in the answer (e.g., “The smart women and rich men” instead of merely “rich men”). We plotted the results for the 3-layer dist probe, the best-performing probe for the Mask model, on both QA corpora in Figure 7. In both plots, the choice of layer in which counterfactuals were inserted had a greater effect than which parse was used to generate the counterfactuals – a sign of poor performance. Depth and other distance probes performed no better.

Visualizing dependency trees for QA embeddings revealed that the counterfactual embeddings induced the correct structure, indicating that the QA model simply did not use such structure in downstream predictions. Furthermore, given the success of our probes and technique with the Mask model, these poor results for the QA model suggest (but admittedly cannot definitely prove) that it may not have learned to use the syntactic information detected by the probes. This theory is consistent with prior art that finds that fine-tuning on specific tasks, as was done for the QA model, worsens the alignment between model and human representations of language (Gauthier and Levy, 2019).

³Plots for the effects of counterfactuals for all probes, models, and test corpora were included in the appendix.

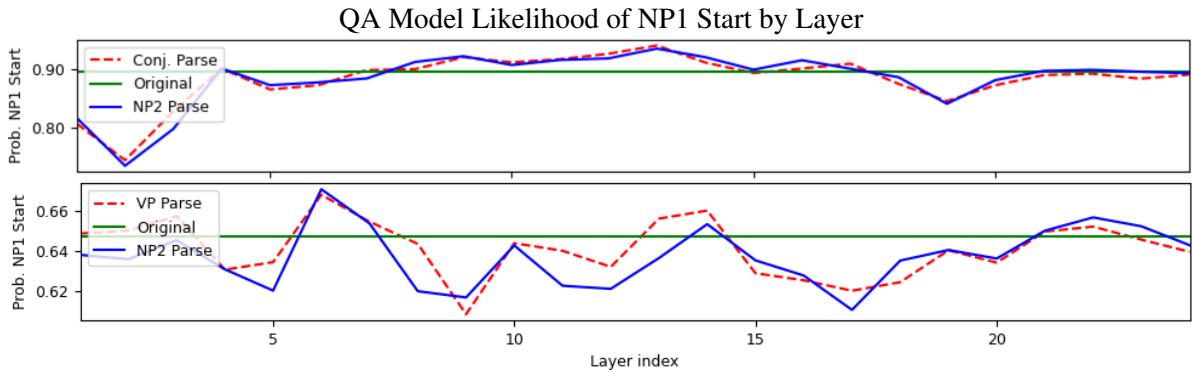


Figure 7: Mean effects of using counterfactual updates from the 3-layer dist probe on the QA model for the RC (top) and NP/VP (bottom) corpora.

6 Conclusion

In this work, we proposed and evaluated a new technique for producing counterfactual embeddings that tested syntactic understanding of models and probes. On the one hand, we uncovered clear evidence supporting a causal link between a distance-based representation of syntax and the outputs of a masked-word model. On the other hand, depth-based manipulations of embeddings had little effect, and we found no evidence that the BERT model finetuned on question-answering uses the syntactic information used by probes.

Our work is merely an initial step in the direction of causal analysis of language models. Developing new probes, backed by causal evidence, could increase our understanding of such models. In particular, our findings that multi-layered probes outperformed linear probes indicate that the prior guidance of simpler probes being preferable may be misleading. Furthermore, as the discrepancy between distance- and depth-based probes revealed, developing a large suite of probe types that focus on different features may be necessary to reveal a model’s reasoning. In tandem with probe development, more sophisticated counterfactual generation techniques than our gradient-based method could produce more interesting counterfactuals for evaluation.

Acknowledgments

We thank the reviewers for their thoughtful comments, in particular regarding adversarial attacks. We thank Professors Julie Shah and Jacob Andreas for ongoing discussions and guidance. Lastly, we thank John Hewitt and Christopher Manning for releasing high-quality, reproducible code, enabling us to rapidly build upon their syntactic probe code-

base.

RPL gratefully acknowledges support from the MIT-IBM Artificial Intelligence Laboratory and MIT’s Quest for Intelligence.

References

- Guillaume Alain and Yoshua Bengio. 2018. [Understanding intermediate layers using linear classifier probes](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Andy Coenen, Emily Reif, Ann Yuan, Been Kim, Adam Pearce, Fernanda Viégas, and Martin Wattenberg. 2019. [Visualizing and measuring the geometry of bert](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8594–8603. Curran Associates, Inc.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single \$\&\!#\ast\$ vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association*

- for *Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2020. [Amnesic Probing: Behavioral Explanation with Amnesic Counterfactuals](#). *arXiv e-prints*, page arXiv:2006.00995.
- Lyn Frazier and Keith Rayner. 1982. Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14:178–210.
- Jon Gauthier and R. Levy. 2019. Linking artificial and human neural representations of language. In *EMNLP/IJCNLP*.
- Yash Goyal, Amir Feder, Uri Shalit, and Been Kim. 2019. Explaining classifiers with causal concept effect (cace). *arXiv preprint arXiv:1907.07165*.
- Rowan Hall Maudslay, Josef Valvoda, Tiago Pimentel, Adina Williams, and Ryan Cotterell. 2020. [A tale of a probe and a parser](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7389–7395, Online. Association for Computational Linguistics.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.
- John Hewitt and Christopher D Manning. 2019. A structural probe for finding syntax in word representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. [Learning the difference that makes a difference with counterfactually-augmented data](#). In *International Conference on Learning Representations*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Comput. Linguist.*, 19(2):313–330.
- Judea Pearl and Dana Mackenzie. 2018. *The Book of Why: The New Science of Cause and Effect*, 1st edition. Basic Books, Inc., USA.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [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. Association for Computational Linguistics.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. [Intriguing properties of neural networks](#). In *International Conference on Learning Representations*.
- Whitney Tabor and Sean Hutchins. 2004. Evidence for self-organized sentence processing: Digging-in effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2):431.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Simas Sakenis, Jason Huang, Yaron Singer, and Stuart Shieber. 2020. [Causal mediation analysis for interpreting neural nlp: The case of gender bias](#).
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.

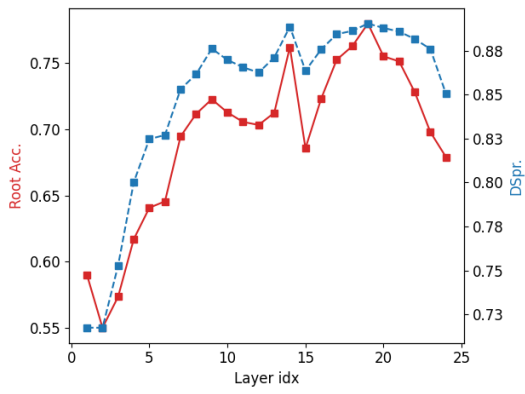
Appendix: Complete Performance Plots

In this appendix, we included additional figures that we were unable to include within the main paper limits.

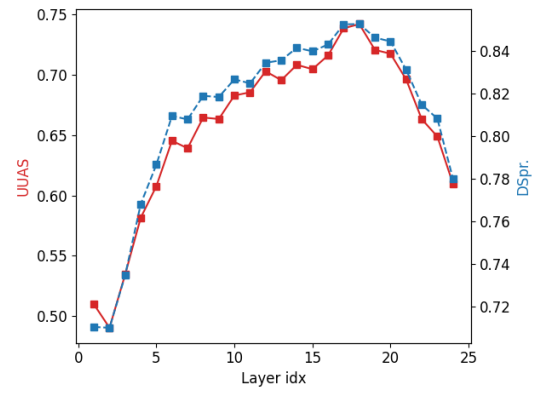
First, we depicted the probe performance characteristics for the 4 probes types we used in all our experiments: the depth, dist, 2-layer dist, and 3-layer dist probes. Each type of probe was trained for both the QA and Mask models. Evaluation of these probes was plotted in Figure 8.

Next, we reported the effect of counterfactual embeddings generated for each model, corpus, and probe type. Given the 4-page limit for the appendix, further plots breaking down the NP/Z corpus, for example, or depicting performance for multi-layered depth probes were not included. These plots merely confirmed trends already present in the data: that depth-based probes did not produce useful counterfactuals, and that the curated and automatically-generated sentences that formed the full NP/Z corpus yielded similar results.

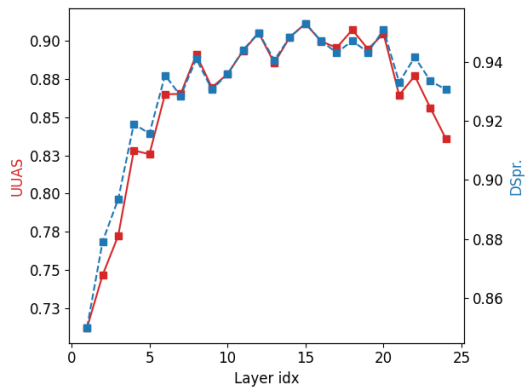
In general, we observed small effects for counterfactuals in the QA Model (Figures 11 and 12), but consistent effects in the Mask Model (Figures 9 and 10). Within the Mask model results, we also observed that the distance probe (2^{nd} row) outperformed the depth probe (1^{st} row), and that the multi-layer distance probes (3^{rd} and 4^{th} rows) outperformed the linear distance probe.



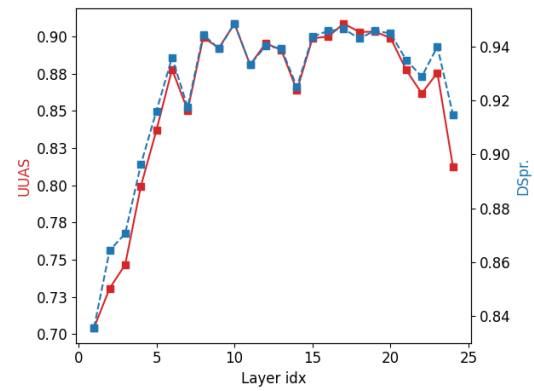
(a) Mask Model Depth Probe



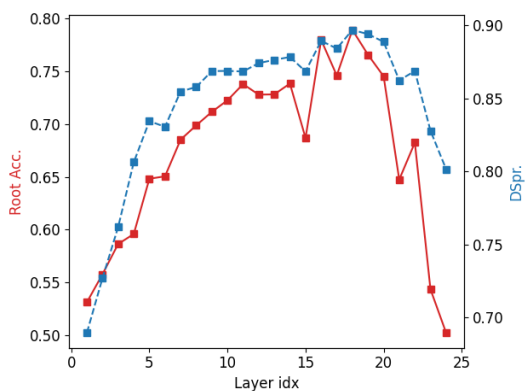
(b) Mask Model Dist Probe



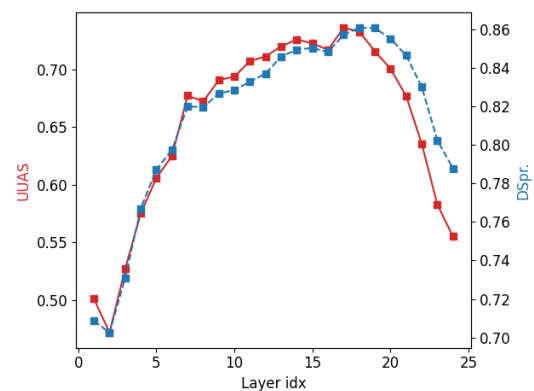
(c) Mask Model 2-Layer Dist Probe



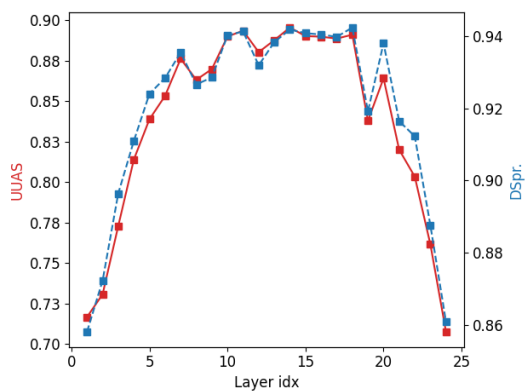
(d) Mask Model 3-Layer Dist Probe



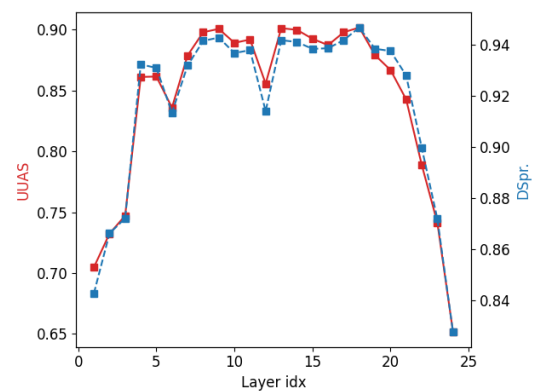
(e) QA Model Depth Probe



(f) QA Model Dist Probe



(g) QA Model 2-Layer Dist Probe



(h) QA Model 3-Layer Dist Probe

Figure 8: Probe performances for the Mask and QA models. Note the changed y axes, demonstrating improved performance for the multi-layer distance probes.

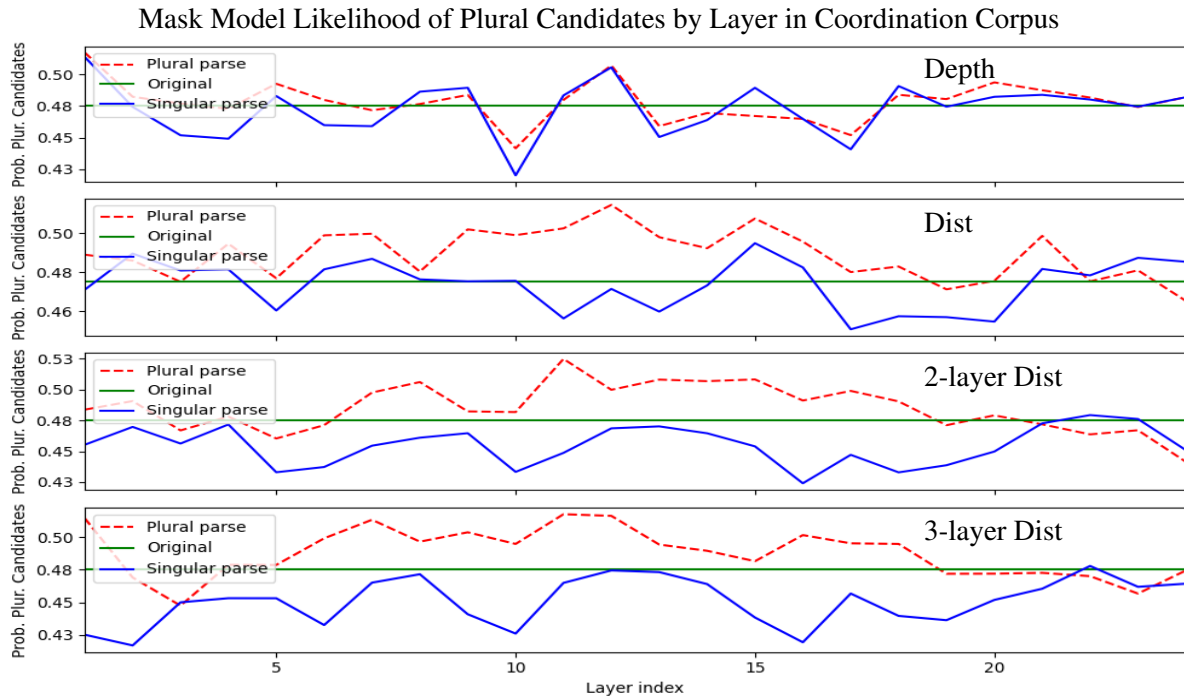


Figure 9: Mask model performance on the Coordination corpus. When using distance-based probes, the plural parse increased the likelihood of plural candidates being predicted, and the singular parse increased the likelihood of singular candidates being predicted.

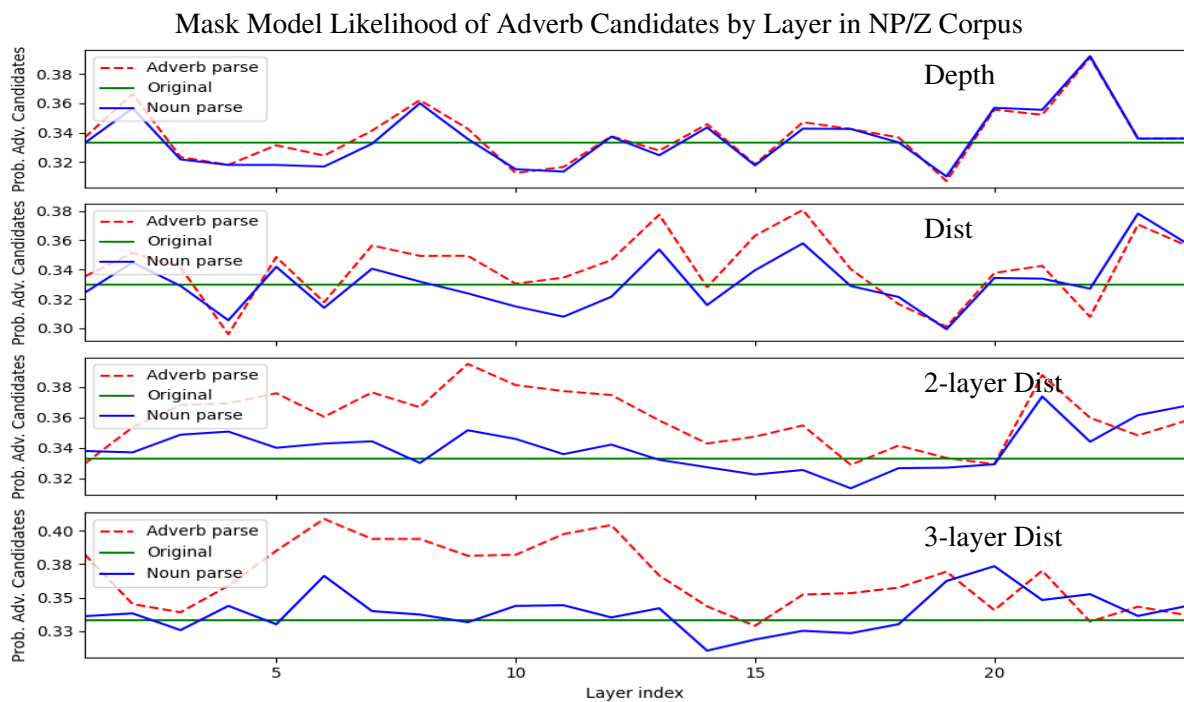


Figure 10: Mask model performance on the NP/Z corpus. Distance-based probes, and in particular multi-layer distance probes, changed model outputs according to syntactic principles.

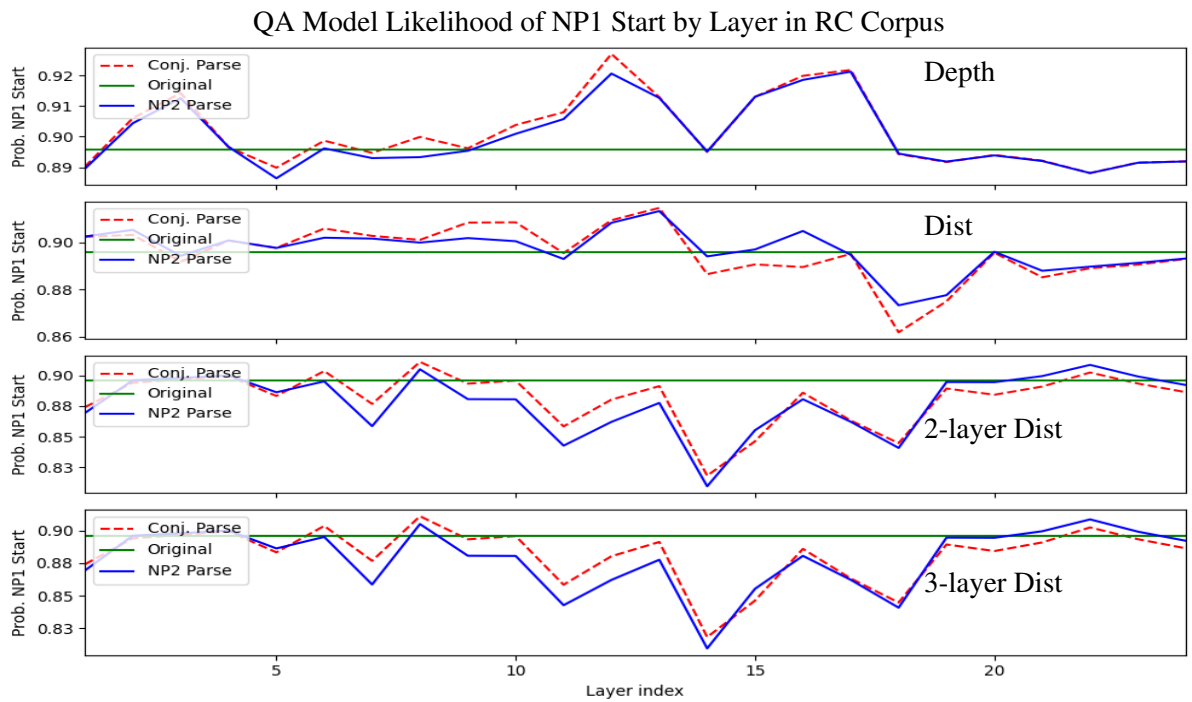


Figure 11: QA model performance on the RC corpus. No probe created consistent effects via counterfactual embeddings.

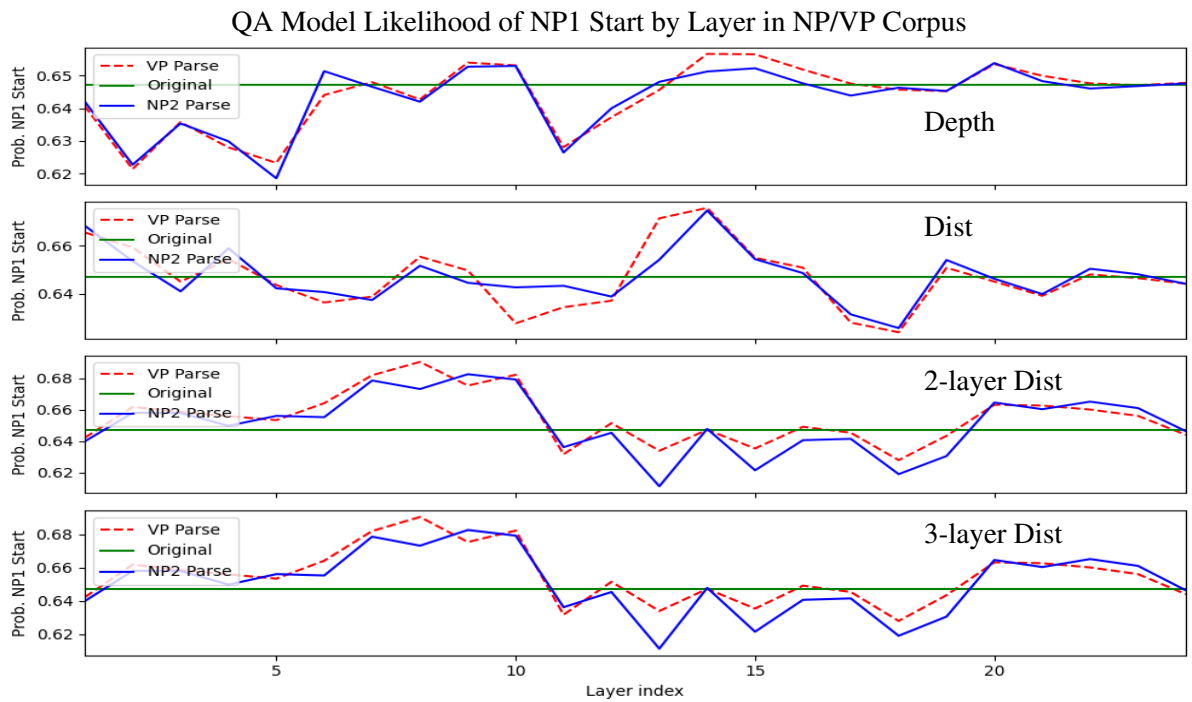


Figure 12: QA model on the NP/VP corpus. As in Figure 11, no probe created consistent effects.