

# “No, they did not”: Dialogue response dynamics in pre-trained language models

Sanghee J. Kim<sup>1</sup>, Lang Yu<sup>2</sup>, Allyson Ettinger<sup>1</sup>

<sup>1</sup>Department of Linguistics, University of Chicago

<sup>2</sup>Meta

{sangheekim, aettinger}@uchicago.edu, langyu@fb.com

## Abstract

A critical component of competence in language is being able to identify relevant components of an utterance and reply appropriately. In this paper we examine the extent of such dialogue response sensitivity in pre-trained language models, conducting a series of experiments with a particular focus on sensitivity to dynamics involving phenomena of at-issueness and ellipsis. We find that models show clear sensitivity to a distinctive role of embedded clauses, and a general preference for responses that target main clause content of prior utterances. However, the results indicate mixed and generally weak trends with respect to capturing the full range of dynamics involved in targeting at-issue versus not-at-issue content. Additionally, models show fundamental limitations in grasp of the dynamics governing ellipsis, and response selections show clear interference from superficial factors that outweigh the influence of principled discourse constraints.

## 1 Introduction

Competence in language involves understanding complex principles governing relevance of previous content and dynamics of referring back to that content. Certain parts of an utterance are more central and more likely to receive a response than others, and the pragmatic and grammatical rules governing responses in dialogue interact with the nature of the content being responded to. Humans are highly sensitive to these distinctions, and we can expect these sensitivities to be critical for robust models in NLP, and especially for dialogue.

Here we examine sensitivity to these dialogue response dynamics in pre-trained language models (PLMs). PLMs are now used as foundation for nearly every downstream NLP task, including dialogue applications (e.g., Upadhye et al., 2020; Koto et al., 2021). The impressive downstream performance enabled by these models has raised important questions about what types of linguistic

competence are being learned during pre-training—and though there is a growing body of work answering aspects of this question, topics of pragmatic and dialogue competence have been relatively understudied. In this paper we focus on addressing this gap, and in particular on understanding the extent to which PLMs develop sensitivity to dynamics governing responses in dialogue. Though these PLMs are not trained to engage in dialogue per se, they can be expected to encounter dialogue during training (in, for instance, novels), so it is not unreasonable to expect that they may learn about such dialogue dynamics along with other linguistic competences. The strength of these models’ sensitivity to such dynamics has important implications for robustness in dialogue applications, since a strong grasp of dialogue dynamics in standard PLMs stands to reduce fine-tuning needs and enable more robust downstream behaviors.

We begin with the notion of *at-issueness*. A component of an utterance is considered *at-issue* if it is part of the “main point” of the utterance—this is to be contrasted with side comments or mentions of background knowledge, which are not the main focus of the sentence. As we lay out in Section 3.1, the distinction between at-issue and not-at-issue content of an utterance is reflected directly in the nature of responses to that utterance. We thus examine models’ preferences for different responses, to assess whether the preferences reflect understanding of at-issueness and how to respond to it. We find that models show consistent preference to target at-issue (main clause) content, but mixed and overall fairly weak sensitivity when it comes to the full range of dynamics involved with at-issueness.

These assessments of at-issueness sensitivity are also critically reliant on another aspect of dialogue response dynamics: ellipsis. We thus additionally make a closer examination of the extent to which constraints from context dictate models’ selection of auxiliary verbs (such as *did*, *does*, *would*) in el-

ipsis constructions. We find that although models often favor an auxiliary verb that targets the main clause, they also make frequent errors, and they very rarely favor both of the auxiliary forms that align with the prior context. These results furthermore raise the important possibility that models are highly sensitive to preferences for particular auxiliary verb types, and that this could drive the at-issueness results as well. With this in mind we revisit the at-issueness experiments, and find that, indeed, there are substantial differences in models’ preferences depending on the identity of the particular verb that targets the relevant content.

Overall, our results suggest that PLMs have non-trivial gaps in their understanding of response dynamics in dialogue. Our results also indicate certain differences between models: BERT and RoBERTa show strong bias toward selecting responses that target the most recent and/or main clause content, while other models show more reliance on individual auxiliary verb properties. In all cases the results indicate that these PLMs have not yet achieved ideal sensitivity to response dynamics involving at-issueness and ellipsis, and that effectiveness in dialogue will benefit from additional training approaches. We make all datasets and code available for further testing.<sup>1</sup>

## 2 Related work

Recent years have seen extensive work on analysis of PLMs. Methodologically, some of the most popular analysis paradigms targeting model embeddings have included classification-based probing (e.g., Kim et al., 2019; Zhang et al., 2019) and correlation with similarity judgments (Finkelstein et al., 2001; Gerz et al., 2016; Conneau and Kiela, 2018). Other work has analyzed PLMs by eliciting and analyzing output predictions (Linzen et al., 2016; Goldberg, 2019). Our work here focuses primarily on the latter methodology, examining and comparing model output probabilities—however, our analysis in Section 5.4 uses classification-based probing. Our work also builds on approaches implementing specialized sentence generation systems that produce large annotated datasets (Ettinger et al., 2018; McCoy et al., 2019).

Analyses of PLMs have targeted a variety of types of linguistic competence. In particular, a large body of work has studied the extent to which

PLMs capture syntactic and semantic information (Linzen et al., 2016; Peters et al., 2018; Bacon and Regier, 2019; Hewitt and Manning, 2019; Tenney et al., 2019). Less work has addressed the extent to which PLMs show sensitivity to pragmatic and discourse information, as we focus on in this paper. Kurfali and Östling (2021) study multilingual models in various discourse tasks via zero-shot learning. Pandia et al. (2021) investigate LMs’ pragmatic competence to predict discourse connectives. Pitler and Nenkova (2009) report that a supervised classifier is able to identify discourse relations given syntactic features along with connectives. Patterson and Kehler (2013) implement a similar idea and show that classifiers are able to predict the presence of a connective based on shallow linguistic cues. Koto et al. (2021) explore pre-trained language models’ capability in capturing discourse level relations. We complement this existing work by branching into new areas of pragmatic and discourse knowledge, examining models’ sensitivity to dialogue response dynamics.

Another closely related literature is that in which PLMs, especially transformer LMs, are used for building dialogue systems directly. Le et al. (2019) propose Multimodal Transformer Networks (MTN) for visual-grounded dialogue tasks. Other work investigates topic-driven language models for emotion detection in dialogues (Zhu et al., 2021). Oluwatobi and Mueller (2020) report state-of-the-art performance on dialogue generation using transformer-based models. There are also language models designed for and trained on dialogue or conversation, such as TransferTransfo (Wolf et al., 2019), PLATO (Bao et al., 2020), ConveRT (Henderson et al., 2020), TOD-BERT (Wu et al., 2020), DialoGPT (Zhang et al., 2020), DialogBERT (Gu et al., 2021), and LaMDA (Thoppilan et al., 2022).

Here we focus on clarifying the extent to which PLMs pre-trained in the standard paradigm can develop knowledge of dialogue dynamics prior to any specialized dialogue training. This line of inquiry serves to broaden our general understanding of linguistic competence of standard PLMs, and also has implications for use of these standard PLMs as foundation for further dialogue-specific training.

## 3 Background

### 3.1 At-issueness

Our analyses focus on the dynamics that govern responses in dialogue, and aspects of prior utter-

<sup>1</sup><https://github.com/sangheek16/dialogue-response-dynamics>

ances that they target. The first notion that we test for in PLMs is sensitivity to “at-issueness.” At-issueness refers to content’s status as the main point of the utterance—to be contrasted with not-at-issue content, such as side comments and assumed knowledge (see [Potts \(2005\)](#) for a comprehensive overview). Humans are sensitive to which content in an utterance is “at-issue” and which content is not—and this sensitivity is reflected in dialogue response dynamics. Consider the utterance in (1).

- (1) The nurse, who has interest in French cuisine, adopted a rescue dog.

If a listener responds to (1) with “No” or “That’s not true,” they would most likely be objecting to the claim that *the nurse adopted a rescue dog*, since this is the main point (at-issue content) of (1). It is less likely that they would be objecting to the side comment about French cuisine. As a result, a response of “No, he didn’t (adopt a rescue dog),” would be natural, while “No, he doesn’t (have interest in French cuisine)” would be less so.

This intuition drives a key diagnostic used to distinguish at-issue and not-at-issue content, known as the **Rejection & Peripherality Test** (or the Assent/Dissent Test) ([Amaral et al., 2007](#); [Koev, 2013](#); [Syrett and Koev, 2015](#)). The “rejection” component of this test is illustrated in (2). Speaker B<sub>1</sub> replies to Speaker A’s utterance with a **rejection** (“No”), and uses the elliptical verb phrase (“did not”) that targets the (at-issue) content of the main clause (“The nurse adopted a rescue dog.”), for a natural and appropriate response. In contrast, Speaker B<sub>2</sub> rejects the (not-at-issue) content inside the appositive relative clause (ARC), which is less natural (indicated with ‘#’).

- (2) a. Speaker A: “The nurse, who has interest in French cuisine, adopted a rescue dog.”  
 b. Speaker B<sub>1</sub>: “No, he did not.” [Targeting at-issue content]  
 c. Speaker B<sub>2</sub>: #“No, he does not.” [Targeting not-at-issue content]

There is, however, a more natural way to object to not-at-issue content: pausing the dialogue to question a side comment or assumption. This is highlighted in the **peripherality test**, which uses phrases like, “Hey, wait a minute” ([von Fintel, 2004](#); [Amaral et al., 2007](#)), or “Wait, this is peripheral to your point but...” ([Koev, 2018](#)) in order

to make targeting not-at-issue content more acceptable. We show an example in (3).

- (3) a. Speaker A: “The nurse, who has interest in French cuisine, adopted a rescue dog.”  
 b. Speaker B: “Wait no, he does not (have interest in French cuisine).” [Targeting not-at-issue content]

Human sensitivity to this pattern of relationship between at-issueness and “No” versus “Wait no” response types has been well attested in psycholinguistic experiments. [Syrett and Koev \(2015\)](#) in their Experiment 1 find that when selecting between responses that target not-at-issue content in an embedded clause of a prior utterance, humans are much more likely to choose a response of type “Wait no” (77%) than of type “No” (23%).<sup>2</sup> By contrast, when selecting between responses that target at-issue content in a main clause of a prior utterance, humans’ rate of selection of these two response types is roughly even. In their Experiment 2, [Syrett and Koev \(2015\)](#) furthermore show that when selecting among “No” type responses, humans have a strong preference for choosing those that target at-issue content of prior utterances (73.9%) compared to not-at-issue content (26.1%).

Leveraging this knowledge of human sensitivities, we make use of diagnostics modeled after the Rejection & Peripherality Test to examine whether PLMs are also sensitive to these discourse dynamics involving at-issueness and response type. For structuring not-at-issue content, we focus on ARCs as used in the examples above.

### 3.2 Ellipsis

The examples above make critical use of the grammatical phenomenon of ellipsis: use of abbreviated verb phrases that refer back to previous verb phrases. In ellipsis, typically an auxiliary verb (like *did*, *does*, *would*) remains as the verb in the elided verb phrase—for instance: “No, he didn’t” is an elided form that could refer back to “The nurse adopted a rescue dog,” standing in for the longer phrase “No, he didn’t adopt a rescue dog.” Ellipsis is another critical component of forming responses in dialogue, and it plays an important prerequisite role in assessing at-issueness. For these reasons, we also test models’ grasp of ellipsis in dialogue.

<sup>2</sup>The specific wordings in this experiment were “Hey, wait a minute,” and “That’s not true.”

## 4 Experiments

### 4.1 Construction of test items

To enable controlled tests inspired by the structure of the Rejection/Peripherality tests, we generate items using templates. Each input item consists of a sequence of two sentences: (a) a **context sentence**, and (b) a **response sentence**.

We generate the **context** sentences based on a core template of “NOUNPHRASE, who VERBPHRASE1, VERBPHRASE2.” This structure includes an embedded ARC (not-at-issue content) and an embedding main clause (at-issue content), as in our example (1) above: *The nurse, who has interest in French cuisine, adopted a rescue dog.* For the noun phrases, we sample from a list of nouns referring to names of occupations (e.g., *nurse, reporter, violinist*). As for verb phrases, to ensure that it would always be unambiguous whether a rejection is targeting the main or the embedded clause, for each item we control the verb phrases of the two clauses such that they will always be targeted by different elided verbs in the response sentence. To do this we create ordered verb pairs from six unique auxiliary verbs *is, was, does, did, has, could*, with the first verb assigned to the embedded clause, and the second assigned to the main clause. This resulted in 30 ( $=_6P_2$ ) unique verb type pairings. We then draw from a list of verb phrases associated with each auxiliary verb: for instance, the verb phrases for the auxiliary verb *does* contain examples such as *has interest in French cuisine, and enjoys hiking*; for the verb *did*, the verb phrases include *adopted a rescue dog, and met the Illinois governor at a Greek restaurant*. We randomly sample from these verb phrases for each verb pair, with the phrase for the first verb assigned to the VERBPHRASE1 position in the template, within the ARC, and the second verb phrase to the VERBPHRASE2 position, in the main clause. Ten unique sentences were generated for each verb pair, resulting in 300 context sentences ( $= 30$  verb pairs \* 10 sentences). Because of our use of the ordered pairs, every auxiliary verb is equally likely to be the correct form for targeting either at-issue content or not-at-issue content of a context sentence.

The **response** sentences then include a “header” consisting of either *No* or *Wait no*, a subject pronoun (sampled randomly to avoid pronoun gender biases), an auxiliary verb targeting either the main clause or embedded clause, and *not*. For example, response sentences might consist of “No, she does

not,” or “Wait no, he has not.” The differences in these headers are the critical factor that impacts whether a response sentence can reasonably target (not-)at-issue content in the context sentence—and the auxiliary verb indicates which verb phrase in the context sentence is being targeted.

In constructing these items, an additional consideration is how to create a setting in which the PLMs may naturally recognize the input as describing a dialogue. We choose to present the items in a format of dialogue resembling that in novels, where entities are described explicitly as uttering the relevant statements. Our final templates thus take a form as shown in example (4) below.<sup>3</sup> We randomly sample the speaker names (e.g., Marco, Ellie) from a list of 400 names, ensuring that no two names repeat in a given item.

- (4) Marco said, “The nurse, who has interest in French cuisine, adopted a rescue dog,” and Ellie replied, “{No / Wait no}, he {did / does} not.”

### 4.2 Models tested

In all of our experiments below, we test six PLMs. Of these models, five are masked language models (MLMs): BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), XLM-RoBERTa (Conneau et al., 2020), DistilBERT and DistilRoBERTa (Sanh et al., 2019). The final model is a causal (unidirectional) language model (CLM): DistilGPT2 (HuggingFace). We used the implementations of these models made available through the HuggingFace Transformers library (Wolf et al., 2020).

## 5 At-issueness tests

### 5.1 Header preferences

We begin by asking whether models, like humans, are sensitive to the role of the response “header” (“No” vs. “Wait no”) in whether a rejection can naturally target (not-)at-issue content. In line with Experiment 1 in Syrett and Koev (2015), we begin by testing whether models recognize that “No” is an appropriate header when the response auxiliary targets the main clause, but “Wait no” is critically more appropriate when the response auxiliary targets the embedded clause. To do this, for a given

<sup>3</sup>We also tested with a simpler dialogue style: A: “The nurse, who has interest in French cuisine, adopted a rescue dog.” B: “{No / Wait no}, he {did / does} not.” This formatting difference did not significantly change the results in the experiments for which we made this comparison.

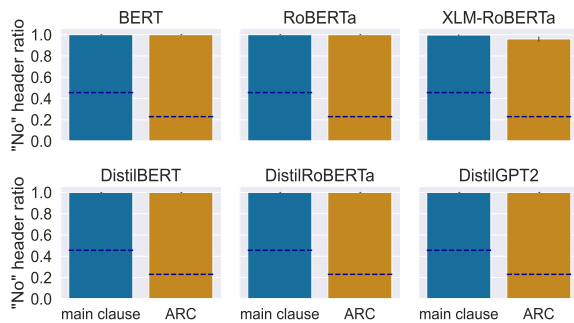


Figure 1: Header selection. The Y-axis indicates the ratio of the “No” header having higher probability than the “Wait no” header given the same verb phrase in the target sentence. The X-axis indicates the target content: ‘main clause’ condition is when the main clause is targeted; ‘ARC’ condition is when the embedded clause is targeted. Dashed lines are human performance baseline reported in Syrett and Koev (2015) (Experiment 1). Error bars = 95% Confidence Interval.

item we hold constant the auxiliary verb in the response sentence (e.g., *did/does/has*), and we compare the model probabilities for headers of “No” vs “Wait no.” The auxiliary verb for a given item either targets the main clause (at-issue) content of the context sentence, or targets the embedded clause (not-at-issue) content. Since these items are different lengths depending on the choice of header, we compare the conditional log probability of the full sequence, normalized by length, for both MLM and CLM models. For MLMs, we compute pseudo-log-likelihoods, which are obtained by summing the conditional log probabilities of each sentence token (as in Salazar et al. (2020)), and normalizing by number of input tokens.<sup>4</sup>

Figure 1 shows the percentage of items for which the model assigns a higher probability to the sequence with “No” than with “Wait no,” separated based on whether the response auxiliary verb targets the main clause (at-issue content) or embedded clause (not-at-issue content). We see that regardless of which clause the response auxiliary targets, models always prefer the “No” header to the “Wait no” header, in a strong contrast with humans’ intuition that “Wait no” is much better for targeting not-at-issue content.

## 5.2 Comparing auxiliary preferences

While the above result casts doubt on models’ sensitivity to the relationship between at-issueness and response headers, we might wonder whether “No”

<sup>4</sup>We use the minicons library (Misra, 2022) for conditional sequence probabilities for all models.

is simply too strong or too frequent a response. Additionally, the fact that “No” and “Wait no” are different in lengths raises questions about whether the probability comparisons are fair in the MLMs (in which use of full sentence probabilities is also a bit atypical).

To explore a different angle on this question, we therefore shift to a direct examination of models’ preferences for response auxiliaries that target the at-issue content, versus those that target not-at-issue content—and how this is affected by the nature of the “No” versus “Wait no” header. This allows us to examine the MLMs in a more natural setting (assessing probabilities on a single masked position), and also allows us to examine how headers impact models’ choices for what contextual content should be targeted.

In these experiments, for MLMs we simply place a [MASK] token at the response auxiliary position and compare auxiliary probabilities at that position:

- (5) Marco said, “The nurse, who has interest in French cuisine, adopted a rescue dog,” and Ellie replied, “{No / Wait no}, he [MASK] not.”

For the CLM, we compare probabilities for the full sequence, with one of two candidate auxiliaries in the target position (as in the previous experiment). The two candidate auxiliaries that we insert are simply the two most relevant: the auxiliary that targets the main clause (at-issue content), and the one that targets the embedded clause (not-at-issue content).

Figure 2 shows the percentage of the time that each model assigns higher probability to the auxiliary targeting the at-issue content, over the auxiliary targeting the not-at-issue content.<sup>5</sup> We see that all models prefer the at-issue-targeting auxiliary at a rate greater than chance, with some models (BERT, RoBERTa) showing preference for targeting the at-issue content almost 100% of the time. The question, then, is whether the use of “Wait no” reduces the rate of targeting the main clause—given that this header allows for targeting of not-at-issue content. A one-sided *t*-test shows that the selection ratio of the at-issue content is indeed larger with the “reject” header compared to the “wait” header

<sup>5</sup>Based on the human performance in Syrett and Koev (2015) (Experiment 2), we could expect a humanlike ratio of selecting at-issue content with the “No” header to be 0.789. This version of the human experiment did not obtain a ratio for the “Wait no” header.




Figure 2: Rejection test. The Y-axis indicates the percentage of instances in which models show preference for responses targeting at-issue content over not-at-issue content. The X-axis indicates type of header in the input sequence: ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header. Error bars = 95% Confidence Interval.

in most of the tested models, at reasonable levels of statistical significance (BERT:  $t = 2.716$ ,  $p = 0.003$ ; RoBERTa:  $t = 1.489$ ,  $p = 0.069$ ; XLM-RoBERTa:  $t = 2.115$ ,  $p = 0.017$ ; DistilBERT:  $t = 0.597$ ,  $p = 0.275$ ; DistilRoBERTa:  $t = 2.056$ ,  $p = 0.02$ ; DistilGPT2:  $t = 1.671$ ,  $p = 0.048$ ). This suggests that at least some of the models may have picked up on some relationship between these headers and targeting of at-issue versus not-at-issue content, though it is also clear that these trends are relatively weak.

### 5.3 Conjunction

The above results show that PLMs exhibit a strong preference for auxiliaries that target the main clause of the context sentence. How should we interpret this preference for targeting the main clause? An immediate question that arises is whether this preference could be due to recency alone: in our items, the verb phrase in the main clause of the context is also always the more recent verb phrase before the response sentence. To investigate this possibility, we modify our items to involve two verb phrases in the context sentence, but with the phrases joined by conjunction (6). This means that both verb phrases are now at-issue, and any preference for one over the other can be attributed to recency.

- (6) Marco said, “The nurse has interest in French cuisine and adopted a rescue dog,” and Ellie replied, “{No / Wait no}, he [MASK] not.”

We again compare model probabilities with each of the two valid candidate auxiliaries. Figure 3 shows the percentage of items for which the models prefer the auxiliary that targets the more recent




Figure 3: Conjunction test. The Y-axis indicates percentage of instances in which models show preference for responses targeting more recent verb phrases over more distant verb phrases. The X-axis indicates type of header in the input sequence: ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header. Dashed lines mark chance level (50%). Error bars = 95% Confidence Interval.

verb phrase in the context sentence. It is clear from these results that the trend toward targeting main clause content in Figure 2 cannot be attributed to recency alone: a majority of models are now hovering around 50% in targeting the most recent phrase, with DistilGPT2 in fact showing a preference to target the more *distant* phrase rather than the more recent one. BERT and RoBERTa, by contrast, do both show some bias to target the more recent verb phrase—however, this trend is substantially weaker than the trend in Figure 2, indicating that although these models do prefer to target more recent content, they also show a preference for targeting main clause content over and above this recency bias.

### 5.4 Probing

The results above suggest that at very least, models are sensitive to the fact that embedded clauses (in this case ARCs) have special status in affecting response dynamics—such that models prefer response auxiliaries that target the main clause in the previous context, over and above effects of recency. In this section we briefly confirm that models are sensitive to the differing status of embedded clause content, through a probing experiment testing whether model representations distinguish embedded clause (not-at-issue content) from main clause (at-issue content). To do this, we extract token embeddings from the models and train a classifier to predict whether these tokens are part of a not-at-issue content or an at-issue content. The task is formulated as 3-class classification: contextualized token embeddings from the last hidden layer are used as input for the classifier, and la-

Model	Accuracy (%)
BERT	99.9
RoBERTa	100
XLM-RoBERTa	99.2
DistilBERT	99.4
DistilGPT2	99.5
DistilRoBERTa	100

Table 1: Probing performance using token embeddings from last hidden layers.

bels are generated based on whether the token is 1) part of at-issue content, 2) part of not-at-issue content or 3) neither.<sup>6</sup> Train/test dataset are randomly split for each model, while keeping tokens from the same input sequence together, yielding on average 8,500 training and 4,000 test samples.<sup>7</sup> For this experiment we use a multi-layer perceptron classifier with a single hidden layer of size 50 with ReLU activation, and a softmax layer to generate binary labels. We use a relatively simple classifier following the reasoning of [Adi et al. \(2017\)](#), that this allows examination of how easily extractable information is in these representations.

As shown in Table 1, all models achieve near perfect classification accuracy. The result further supports the conclusion that these models do encode distinctions between content in the main clauses of these sentences and content in embedded clauses—such that the trends in favor of targeting main clause content may be considered to reflect some real sensitivity to contributions of these structural properties to dialogue dynamics.

## 6 Ellipsis

The response tests above rely on a critical prerequisite: that models understand how to use ellipsis structures like “he didn’t” and “she doesn’t.” In the case of our items, it is specifically the case that there are only two auxiliary verbs that could possibly be appropriate in a given response sentence, because there are only two verb phrases in the context sentence that could be rejected. In this section we take a closer look at whether models’ preferences for response auxiliaries reflect these broader discourse constraints on ellipsis.

<sup>6</sup>Tokens counted as “neither” are those like “Marco said” that are used to introduce the dialogue content.

<sup>7</sup>To mitigate impacts of random variation in train/test split across models, we trained the probe for each model three times and averaged the results.

### 6.1 Ellipsis top one accuracy

We begin by examining the auxiliaries that receive top probability from the models, among the six tested auxiliary verb candidates (i.e., *did*, *does*, *has*, *is*, *was*, and *would*). Specifically, we ask whether the highest-probability response auxiliary selected by the model for a given context is appropriate given the context sentence and header. This test differs from our comparisons above because the previous tests simply compared the two relevant auxiliaries (main clause and embedded clause), without testing whether either of these auxiliaries was assigned the highest probability among all possible auxiliary verbs. Here we count the model as correct if in the case of the “No” header it assigns the highest probability to the auxiliary that targets the main clause, or if in the case of the “Wait no” header it assigns the highest probability to either the main-clause-targeting or embedded-clause-targeting auxiliary (because “Wait no” could also reasonably target the at-issue content).

Figure 4 shows the percentage of the time that the top-ranked auxiliary is among those counted as correct based on the header. We see wide variation in the models’ performance on this assessment, with BERT and RoBERTa preferring the correct auxiliary nearly 100% of the time, but distilled models rarely selecting the correct auxiliary as top choice. This suggests at first glance that BERT and RoBERTa have gained a stronger grasp on the relationship of elided auxiliary forms to the previous context—however, it must also be noted that our definition of “correct” favors BERT and RoBERTa because preference for targeting the main clause (which these two models have exhibited) can always be counted as correct. We thus implement a more difficult ellipsis test in the next section.

### 6.2 Ellipsis top two accuracy

As we describe above, because there are only two verb phrases in each of our context sentences, it is clear that there are only two acceptable auxiliary verb forms that can occur in a given response sentence. To test whether models have a grasp of this constraint, in this section we examine the top two highest-probability auxiliaries, and assess the percentage of the time that these auxiliaries are exactly the two that target the main clause and embedded clause of the context sentence, respectively.

Figure 5 shows the percentage of the time that the two acceptable auxiliaries are the top two




Figure 4: Ellipsis test (top-1 accuracy). The Y-axis indicates the percentage of instances in which models prefer responses targeting the main clause (in the case of the “No” header), or targeting either the main clause or embedded clause (in the case of the “Wait no” header). The X-axis indicates type of header in the input sequence: ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header. Error bars = 95% Confidence Interval.




Figure 5: Ellipsis test (top-2 accuracy). The Y-axis indicates the percentage of instances in which models’ top two highest-probability auxiliary verbs are the two acceptable auxiliary verbs given context. The X-axis indicates type of header in the input sequence: ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header. Error bars = 95% Confidence Interval.

highest-probability auxiliaries for the models. It is clear that the accuracies here are very low—even the most accurate models meet the criterion only 20-30% of the time, suggesting that this category of grammatical/discourse sensitivity is still largely missing from these models.

### 6.3 Error analysis

To get a better sense of where the models are going wrong in these tests, we perform error analyses for both of the two ellipsis tests. For the top-1 ellipsis test we examine cases where the top auxiliary is not “correct,” and for the top-2 test we examine cases where at least one inappropriate auxiliary “intrudes” in the model’s top two. In Figures 7 and 8 in the

Appendix, we show the distribution of auxiliary verbs that the models prefer among these erroneous cases. We see in these figures particularly substantial interference from more frequent auxiliaries like *did*, *does*, and *is*, suggesting that rather than guiding auxiliary choice based primarily on discourse constraints, the model probability distributions are non-trivially influenced by general frequency of the individual auxiliary verbs in ellipsis.<sup>8</sup>

## 7 Verb analysis in rejection test

The previous section raises questions about the extent to which these PLMs have a grasp on the basic discourse constraints that govern ellipsis in response utterances—and this ellipsis serves as critical foundation for our at-issuence response tests in Section 5.2. In particular, the error analysis above indicates that model probabilities are influenced in large part by biases in favor of particular frequent auxiliary verbs. In this section we thus return to our at-issuence test to examine behaviors of individual auxiliary verbs separately.

Figure 6 shows the percentage of the time that the models prefer targeting at-issue content, broken down by which auxiliary verb targets the ARC (top row) or the main clause (bottom row). We see that for four of the models, the identity of the auxiliary verb makes a substantial difference: for instance, distilled models strongly prefer to target the main clause if “did” or “does” is the auxiliary that targets the main clause—but if “did” or “does” targets the embedded clause, these models strongly prefer targeting the embedded clause. In other words, the distilled models appear in large part simply to be biased toward preferring “did” or “does” in the response sentence. BERT and RoBERTa do not show as much verb-specific fluctuation, instead preferring to target the main clause content regardless of which auxiliary verb does so. This suggests that these two models have a relatively more robust grasp on use of ellipsis to target particular clauses in previous context. As for sensitivity to impact of headers on response dynamics, within individual verbs we occasionally see a trend such that the “wait” header results in less targeting of the main clause (XLM-RoBERTa with “has” and “would,” DistilGPT with “was” and “would,” etc.), but for many verbs we see no difference, or even the op-

<sup>8</sup>Among the possible verb phrase ellipsis triggers (e.g., *be*, *has*, *do*, etc.), the auxiliary verb *do* has been reported to be the most frequent (44%), with the auxiliary verb *be* following the next (22%) (Bos and Spénader, 2011).






Figure 6: Verb analysis on rejection test. The Y-axis indicates the percentage of instances in which models show preference for responses targeting at-issue content over not-at-issue content. The X-axis indicates the identity of the response auxiliary that can be used to target the selected clause of the context sentence: in the top row the X-axis indicates the auxiliary that targets content inside the ARC (not-at-issue content) of the context sentence, and in the bottom row the X-axis indicates the auxiliary that targets content inside the main clause (at-issue content) of the context sentence. ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header. Error bars = 95% Confidence Interval.

posite trend. On the whole, the impact of auxiliary verb identity is for most models much stronger than that of header.

## 8 Discussion

In this paper we have reported on a series of experiments testing sensitivity of pre-trained language models to dynamics involved in responding to an utterance in dialogue. We focus specifically on at-issueness and ellipsis, and find that models show clear sensitivity to the special status of embedded clauses, and general preference to target main clause content—but they show mixed results in terms of understanding the interaction of response headers with targeting of at-issue versus not-at-issue content. Furthermore, they show certain basic limitations in their grasp of the principles governing ellipsis, with selection of auxiliaries often influenced by superficial frequency factors rather than principled discourse constraints. Our findings also highlight differences between models, with certain models showing strong preference to target main clause content, and others showing stronger fluctuations based on individual auxiliary verbs.

This work highlights potential for improvement in standard PLMs, with respect to discourse sensitivities that have real implications for language competence generally, and for dialogue in particular. The models’ sensitivity to special status of embedded clauses is consistent with work indicating sensitivity to syntax in these models (Goldberg, 2019), and the consistency with which BERT and

RoBERTa prefer auxiliaries targeting main clause content indicates that these models pick up on some interaction between ellipsis and syntax of previous context. Additionally, the slight impacts of header in Section 5.2 suggest that these models may pick up on the beginnings of a relationship between response types and the types of content that they target. However, the general weakness in sensitivity to headers, failure on many aspects of the ellipsis tests, and interference of superficial factors, indicate clear room for growth in capturing the full range of these discourse dynamics.

From a perspective of downstream dialogue tasks, our findings indicate that discourse competence in standard PLMs is not sufficiently comprehensive to expect that these models can provide a fully robust foundation for dialogue applications. It is possible—though not guaranteed—that training or fine-tuning directly for dialogue could improve the robustness of models’ sensitivity to the specific types of response dynamics tested for here. We leave this question for future work.

## Acknowledgments

We are grateful to Ming Xiang, Shane Steinert-Threlkeld, Tal Linzen, Najoung Kim, and members of the UChicago CompLing Lab, for valuable comments and discussion. We thank Kanishka Misra in particular for guidance on usage of the minicons library. We also thank three anonymous reviewers for their thoughtful feedback and suggestions.

## References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. [Fine-grained analysis of sentence embeddings using auxiliary prediction tasks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Patricia Amaral, Craige Roberts, and E. Allyn Smith. 2007. [Review of the logic of conventional implicatures by Chris Potts](#). *Linguistics and Philosophy*, 30(6):707–749.
- Geoff Bacon and Terry Regier. 2019. [Does bert agree? Evaluating knowledge of structure dependence through agreement relations](#). *arXiv preprint arXiv:1908.09892*.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. [PLATO: Pre-trained dialogue generation model with discrete latent variable](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 85–96, Online. Association for Computational Linguistics.
- Johan Bos and Jennifer Spenser. 2011. [An annotated corpus for the analysis of VP ellipsis](#). *Language Resources and Evaluation*, 45(4):463–494.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau and Douwe Kiela. 2018. [SentEval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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.
- Allison Ettinger, Ahmed Elgohary, Colin Phillips, and Philip Resnik. 2018. [Assessing composition in sentence vector representations](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1790–1801, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. [Placing search in context: The concept revisited](#). In *Proceedings of the 10th International Conference on World Wide Web, WWW '01*, page 406–414, New York, NY, USA. Association for Computing Machinery.
- Daniela Gerz, Ivan Vulić, Felix Hill, Roi Reichart, and Anna Korhonen. 2016. [SimVerb-3500: A large-scale evaluation set of verb similarity](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2173–2182, Austin, Texas. Association for Computational Linguistics.
- Yoav Goldberg. 2019. [Assessing BERT’s syntactic abilities](#). *arXiv preprint arXiv:1901.05287*.
- Xiaodong Gu, Kang Min Yoo, and Jung-Woo Ha. 2021. [DialogBERT: Discourse-aware response generation via learning to recover and rank utterances](#). *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, 35(14):12911–12919.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. [ConveRT: Efficient and accurate conversational representations from transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174, Online. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- HuggingFace. [DistilGPT2](#).
- Najoung Kim, Roma Patel, Adam Poliak, Patrick Xia, Alex Wang, Tom McCoy, Ian Tenney, Alexis Ross, Tal Linzen, Benjamin Van Durme, Samuel R. Bowman, and Ellie Pavlick. 2019. [Probing what different NLP tasks teach machines about function word comprehension](#). In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (\*SEM 2019)*, pages 235–249, Minneapolis, Minnesota. Association for Computational Linguistics.
- Todor Koev. 2013. [Apposition and the structure of discourse](#). PhD dissertation, Rutgers University – New Brunswick.
- Todor Koev. 2018. [Notions of at-issueness](#). *Language and Linguistics Compass*, 12(12):1–16.
- Fajri Koto, Jey Han Lau, and Timothy Baldwin. 2021. [Discourse probing of pretrained language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3849–3864, Online. Association for Computational Linguistics.

- Murathan Kurfalı and Robert Östling. 2021. [Probing multilingual language models for discourse](#). In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepLANLP-2021)*, pages 8–19, Online. Association for Computational Linguistics.
- Hung Le, Doyen Sahoo, Nancy Chen, and Steven Hoi. 2019. [Multimodal transformer networks for end-to-end video-grounded dialogue systems](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5612–5623, Florence, Italy. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Kanishka Misra. 2022. [minicons: Enabling flexible behavioral and representational analyses of transformer language models](#). *arXiv preprint arXiv:2203.13112*.
- Olabiyi Oluwatobi and Erik Mueller. 2020. [DLGNet: A transformer-based model for dialogue response generation](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 54–62, Online. Association for Computational Linguistics.
- Lalchand Pandia, Yan Cong, and Allyson Ettinger. 2021. [Pragmatic competence of pre-trained language models through the lens of discourse connectives](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 367–379, Online. Association for Computational Linguistics.
- Gary Patterson and Andrew Kehler. 2013. [Predicting the presence of discourse connectives](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 914–923, Seattle, Washington, USA. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Luke Zettlemoyer, and Wen-tau Yih. 2018. [Dissecting contextual word embeddings: Architecture and representation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509, Brussels, Belgium. Association for Computational Linguistics.
- Emily Pitler and Ani Nenkova. 2009. [Using syntax to disambiguate explicit discourse connectives in text](#). In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 13–16, Suntec, Singapore. Association for Computational Linguistics.
- Christopher Potts. 2005. *The logic of conventional implicatures*. Oxford University Press.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Kristen Syrett and Todor Koev. 2015. [Experimental evidence for the truth conditional contribution and shifting information status of appositives](#). *Journal of Semantics*, 32(3):525–577.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. [LaMBDA: Language models for dialog applications](#). *arXiv preprint arXiv:2201.08239*.
- Shiva Upadhye, Leon Bergen, and Andrew Kehler. 2020. [Predicting reference: What do language models learn about discourse models?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 977–982, Online. Association for Computational Linguistics.
- Kai von Fintel. 2004. [Would you believe it? the King of France is back! \(Presuppositions and truth value intuitions\)](#). In Anne Bezuidenhout and Marga Reimer, editors, *Descriptions and beyond: An interdisciplinary collection of essays on definite and indefinite descriptions and other related phenomena*, pages 315–341. Oxford: Oxford University Press.
- 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.

Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. **TransferTransfo: A transfer learning approach for neural network based conversational agents**. *arXiv preprint arXiv:1901.08149*.

Chien-Sheng Wu, Steven C.H. Hoi, Richard Socher, and Caiming Xiong. 2020. **TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929, Online. Association for Computational Linguistics.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. **DIALOGPT : Large-scale generative pre-training for conversational response generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.

Yuan Zhang, Jason Baldridge, and Luheng He. 2019. **PAWS: Paraphrase adversaries from word scrambling**. 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.

Lixing Zhu, Gabriele Pergola, Lin Gui, Deyu Zhou, and Yulan He. 2021. **Topic-driven and knowledge-aware transformer for dialogue emotion detection**. In *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)*, pages 1571–1582, Online. Association for Computational Linguistics.

## A Appendix: Error analysis on ellipsis and auxiliary preference




Figure 7: Verb analysis on ellipsis test (top-1). The Y-axis indicates the proportion of the verb appearing as the top-1 prediction even when the corresponding auxiliary verb did not appear in the input sequence. Proportion is calculated by header. The X-axis shows the auxiliary verb used in the target sentence. ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header.




Figure 8: Verb analysis on ellipsis test (top-2). The Y-axis indicates the proportion of the verb appearing as one of top-2 predictions even when the corresponding auxiliary verb did not appear in the input sequence. Proportion is calculated by header. The X-axis shows the auxiliary verb used in the target sentence. ‘Reject’ = “No” header; ‘Wait’ = “Wait no” header.