

Do Language Models Know Theo Has a Wife? Investigating the Proviso Problem

Tara Azin^{1,*,\dagger}, Daniel Dumitrescu^{2,*}, Diana Inkpen², Raj Singh¹

¹Carleton University, Canada

²University of Ottawa, Canada

^{\dagger}Corresponding author: taraazin@cmail.carleton.ca

*Equal contribution

Abstract

We investigate how language models handle the proviso problem, an unresolved issue in pragmatics where presuppositions in conditional sentences diverge between theoretical and human interpretations. We reformulate this phenomenon as a Natural Language Inference task and introduce a diagnostic dataset designed to probe presupposition projection in conditionals. We evaluate RoBERTa, DeBERTa, LLaMA, and Gemma using explainability analyses. The results show that models broadly align with human judgments but rely on shallow pattern matching rather than semantic or pragmatic reasoning. Our work provides the first computational evaluation framework for the proviso problem and highlights the need for diagnostic, multi-method approaches to assess pragmatic competence and context-dependent meaning in language models.

Keywords: presupposition projection, proviso problem, pragmatic reasoning, language model evaluation, explainability

1. Introduction

Pragmatics, or the study of context in linguistics, remains an underexplored area in the evaluation of Large Language Models (LLMs). One core pragmatic phenomenon is presupposition, where part of an utterance is assumed to be true and taken as shared knowledge between interlocutors (Strawson, 1950; Stalnaker, 1973, 1998). For example, in the sentence *Theo stopped smoking*, the common ground shared by speaker and listener is that Theo used to smoke (the presupposition). Presuppositions are typically signaled by specific lexical items or constructions known as triggers (e.g., stop, again, possessive pronouns).

While straightforward in simple cases, presupposition becomes more complex in embedded structures, such as in conditional sentences. According to formal semantic theories (Karttunen and Peters, 1979; Heim, 1992; Geurts, 1996), in conditionals of the form $S = \text{If } A, B_p$, where p is the direct presupposition of the consequent B , the presupposition of the sentence is predicted to project conditionally¹. For example, in the sentence *If Theo hates sonnets, so does his wife*, the predicted presupposition is *If Theo hates sonnets, then Theo has a wife*.

However, in practice, speakers typically accommodate a stronger, unconditional presupposition: *Theo has a wife*. This phenomenon in pragmatics and formal semantics is called the proviso problem

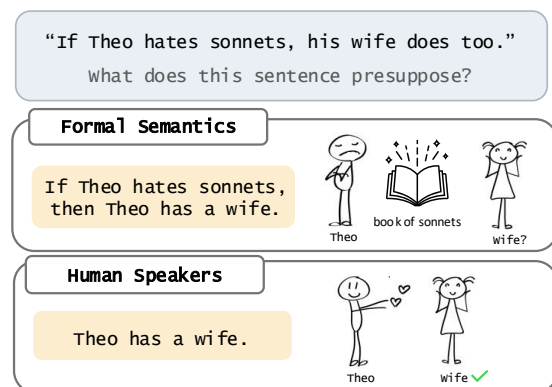


Figure 1: The proviso problem illustrated. Given the sentence *If Theo hates sonnets, so does his wife*, formal semantic theories predict a conditional presupposition, while human speakers typically accommodate an unconditional presupposition. Our work investigates where language models fall in this theory-human divide.

(Geurts, 1996), which remains unresolved². This discrepancy between what *should be* presupposed and what speakers *prefer* to presuppose motivates our investigation of how language models handle presuppositions in conditional structures (Figure 1).

To the best of our knowledge, our work is the

¹Projection here refers to the persistence of presuppositions under embedding in larger structures, such as negation or conditionals (Karttunen, 1973; Heim, 1983).

²The problem is often explained by presupposition accommodation, where the listener updates the context to satisfy the sentence's projected semantic presupposition (Lewis, 1979; von Stechow, 2008; Singh, 2020).

first systematic evaluation of how language models address the proviso problem. We examine whether language models project presuppositions in conditionals as predicted by formal semantic theories or in ways that align with human judgments. To do this, we formulate the proviso problem as a Natural Language Inference (NLI) task. NLI is a well-established framework for modeling inferential relations between a premise and a hypothesis with labels such as Entailment, Neutral, and Contradiction (Bowman et al., 2015; Williams et al., 2018). In our formulation, the premise takes the form $S = \text{If } A, B_p$, where p is the presupposition of B , and the hypothesis is p itself. We then assign NLI labels: Entailment (E) when p follows as the presupposition of S , or Neutral (N) when no inferential relationship holds. For example, given the premise *If Theo hates sonnets, so does his wife* and hypothesis *Theo has a wife*, the theory-based label would be N, as the conditional presupposition does not unconditionally entail the hypothesis. In contrast, humans assign E, reflecting the unconditional presupposition that speakers typically accommodate. The contrast between theory-based predictions and human judgments in our evaluation is intended as a methodological design choice for probing model behavior, rather than as a claim that existing theoretical accounts of presupposition projection are inadequate.

For evaluation, we use explainability techniques, such as gradient-based saliency and attention analysis (Sundararajan et al., 2017; Atanasova et al., 2020), to examine how models come up with their judgments regarding the relation between conditional sentences and their presuppositions. While accuracy metrics show whether a model selects the correct label, they do not make clear whether the decision is driven by sensitivity to presupposition triggers or by reliance on superficial patterns in the data. Our evaluation method addresses this gap by tracking which lexical items and structural features influence model predictions, helping determine whether model behavior aligns with pragmatic reasoning or instead shows reliance on surface-level heuristics.

Our main contributions are as follows:

- We reformulate the proviso problem as an NLI task to make presupposition projection computationally testable.
- We create the first dataset for the proviso problem with approximately 8,500 examples featuring structural, semantic, and contextual variations.
- We evaluate neural language models against theoretical predictions and human judgments using saliency methods to analyze their presuppositional reasoning.

2. Related Work

The evaluation of presuppositional reasoning in language models is a relatively recent area of investigation. Early work has examined logical inference in LLMs, including reasoning about conditionals (Holliday et al., 2024) and presupposition judgment (Atwell et al., 2025), primarily through prompt-based approaches without model fine-tuning. Most prior work on presupposition evaluation has framed the task in the NLI format, creating datasets designed for general inference rather than specifically targeting presuppositional phenomena. Datasets such as IMPPRES (Jeretic et al., 2020) and NOPE (Parrish et al., 2021) include some presuppositional content, but they either focus on simple conditional forms or lack coverage of structurally complex cases such as embedded conditionals.

More recently, CONFER (Azin et al., 2025) introduced a dedicated NLI benchmark for evaluating presuppositional reasoning in conditionals. The dataset includes five types of conditional sentences with human annotations based on the relationship between premises and hypotheses. Their results demonstrate that fine-tuning on this data improves model performance on presuppositional reasoning tasks. However, existing evaluation methods rely primarily on classification accuracy and do not examine whether models process presupposition triggers in ways that align with human reasoning or whether structural variations affect presupposition projection.

Building on this foundation, our work addresses this methodological gap by combining classification-based evaluation with attribution-driven saliency and attention analyses to examine how models process presuppositions in conditional contexts.

3. Dataset Construction

We built our dataset on top of CONFER (Azin et al., 2025), which contains conditional sentences paired with their presuppositions. From this dataset, we selected 900 sentence pairs. In each pair, the premise follows the form $S = \text{If } A, B_p$, where p represents the presupposition of B , and the hypothesis restates p itself. Each selected pair is annotated as an entailment relation between the premise and the hypothesis. The conditional sentences contain presupposition triggers *again* and possessive pronouns (e.g., *her*, *him*).

In this initial subset, half of the examples show a logical dependency between the antecedent A and the hypothesis p . We refer to these as *dependent* cases (DEP), where the content of A semantically relates to or contextually supports p . In the other half, A and p are logically unrelated, which we label as *independent* cases (IND), where A describes

an unrelated event or circumstance. Distinguishing between DEP and IND examples follows existing dynamic semantics theories, which suggest that the relation between the components of conditionals directly affects whether listeners interpret the presupposition conditionally or unconditionally (Singh, 2007; von Stechow, 2008). Examples are shown in Table 1.

We operationalize DEP versus IND as an intra-sentential licensing relation between the antecedent A and the presupposition p of the consequent B : DEP if A plausibly supports p , IND if it does not. Because our materials are single-sentence pairs, this notion of support is restricted to what is inferable within the sentence.

Dependent case (DEP)	
Premise	<i>If Randolph is a carpenter, he'll use <u>his</u> beading tools for designing.</i>
Hypothesis	<i>Randolph has beading tools.</i>
Independent case (IND)	
Premise	<i>If Lisa finishes the meeting early, she'll never drive a sports car <u>again</u>.</i>
Hypothesis	<i>Lisa has driven a sports car before.</i>

Table 1: Examples of DEP and IND cases. In DEP, the antecedent A logically relates to the presupposition p (here, being a carpenter relates to using beading tools). In IND, A and p are unrelated (finishing a meeting has no connection to driving sports cars). Presupposition triggers are underlined.

We expanded the dataset to approximately 8,500 sentence pairs by modifying the original 900 examples. The expansions focused on both structural and semantic properties of the conditionals, including syntactic variation, changes in trigger-hypothesis relatedness, and premise-level contextual modifications. The original labels were largely retained as E , with a small subset reclassified as N . In addition, we added a new column containing theory-based labels, which are consistently marked as neutral for all sentence pairs. This is based on the theoretical accounts of the proviso problem that accept only the conditional form of presupposition, in contrast to the unconditional hypotheses represented in our dataset.

The entailment labels used in our dataset follow the annotation scheme of the CONFER dataset (Azin et al., 2025), on which our dataset is built. For the original 900 sentence pairs, we retain the labels provided in CONFER. During dataset expansion, modified examples were created and reviewed by a computational linguist to ensure that the semantic relation between the premise and the hypothesis remained consistent with the original annotation scheme. Because the modifications were designed to preserve the original entailment relation, the la-

els were largely retained as E , with a small subset reclassified as N when the modification affected the inference relation. Since the annotation and verification were conducted by a single annotator, inter-annotator agreement does not apply.

The additional theory-based label column was generated deterministically based on theoretical accounts of the proviso problem (Singh, 2007; von Stechow, 2008). Under these accounts, presuppositions project conditionally from conditional antecedents, which means that unconditional hypotheses of the form used in our dataset are predicted to be neutral with respect to the premise. Consequently, all sentence pairs receive a neutral theory-based label. The resulting dataset is organized into four subsets.³ Table 2 provides examples from each subset. In this section, we describe the structure of each subset in more detail.

Naming conventions. Throughout this paper, we use the following convention to refer to specific data types. DEP conditionals include only possessive triggers (e.g., *his wife*), whereas IND conditionals include either possessive triggers or the adverbial trigger *again*. Accordingly, each IND example carries either the $-poss$ (for possessive) or $-again$ suffix. We further distinguish between the original examples drawn directly from CONFER (e.g., *DEP*, *IND-poss*, *IND-again*) and their modified variants created for our experiments (e.g., *DEP-mod*, *IND-poss-mod*, *IND-again-mod*). When no $-mod$ suffix is present, the example refers to the original, unmodified instance from CONFER.

3.1. Subset 1: Original Sentences

This subset used the initial 900 NLI pairs without modifications as a baseline to evaluate models' performance.

3.2. Subset 2: Structural Variation

To test how structural manipulation affects presupposition projection, we created three modified versions of each original sentence: (1) Conjunction: adding a conjunct to the antecedent (e.g., *If A and B, then C*); (2) Disjunction: rephrasing the conditional as *Either not-A or B*; and (3) Belief Embedding: embedding the consequent under an attitude verb (e.g., *X believes that B*).

These modifications are based on established findings in presupposition theory showing that presuppositions behave differently across various sentence structures (Karttunen, 1973; Heim, 1983). In conjunctions, presuppositions from both parts typically survive. In disjunctions, presuppositions

³The dataset and experimental results are available at <https://github.com/Conditional-NLI/PROVISER>.

may be blocked or filtered out. Finally, under belief verbs, presuppositions may or may not project depending on the specific verb used.

3.3. Subset 3: Trigger–Hypothesis Relatedness

In this subset, we modified the trigger phrase by manipulating its semantic similarity to the hypothesis. The original examples naturally contained triggers that were semantically *related* to their hypotheses. Using WordNet (Miller, 1995) and ConceptNet (Speer et al., 2017) relations (e.g., *is-a*, *part-of*, *hyponymy*), we generated two additional variants, *somewhat related* and *unrelated*, by substituting key lexical items in the presupposition trigger, thus creating three levels of semantic relatedness for each example.

For example, consider the original premise *If Matt is a scuba diver, he'll bring his wetsuit* with the hypothesis *Matt has a wetsuit*. Here are examples with two degrees of relatedness between the hypothesis and the trigger phrase:

- **Somewhat related:** If Matt is a scuba diver, he'll bring his garment. (garment is a hypernym of wetsuit)
- **Unrelated:** If Matt is a scuba diver, John's friend will bring her wetsuit. (mismatched agent and possessor)

**Hypothesis for both: Matt has a wetsuit.*

As the semantic relationship between the trigger and hypothesis weakens, the gold labels were changed to show the reduced support for entailment. In this example, both the *somewhat related* and *unrelated* variants received N labels instead of the original E, since the premise no longer provides clear evidence for the hypothesis.

3.4. Subset 4: Context–Trigger Relatedness

This subset examines how the semantic relationship between different parts of the conditional affects presupposition projection. We modified the logical connection between the antecedent and consequent to be either logically related or unrelated.⁴

Each sentence contains a presupposition trigger embedded within a conditional structure. We created the two variants by altering the semantic relationship between the conditional's antecedent and consequent. This approach is motivated by

⁴The logical relation between *A* and *p* remained the same, therefore we still have the same DEP and IND sentence types.

research showing that presuppositions can be sensitive to the broader discourse context and the logical coherence of the surrounding material (Geurts, 1999; Beaver, 2001).

The following examples illustrate these two variants:

- **Related:** If Sarah attends a movie festival, she'll never watch Star Wars again.
- **Unrelated:** If Sarah attends the conference, she'll never watch Star Wars again.

**Hypothesis for both: Sarah has watched Star Wars before.*

To facilitate attention analysis (detailed in Section 4.2), we annotated two key regions in each example: K1, the noun phrase preceding the presupposition trigger in the hypothesis (e.g., *Star Wars* in the examples above), and K2, the contextual phrase in the consequent that provides semantic context (e.g., *attends a movie festival*). We use these marked regions to analyze how models link the trigger to its surrounding context.

According to presupposition theory (Karttunen, 1973; Heim, 1983), the presupposed content should project regardless of the logical relationship between the antecedent and consequent, since presuppositions are typically preserved across different contexts. However, we hypothesize that language models may be sensitive to these contextual manipulations, potentially showing different labeling patterns when the conditional's parts are semantically unrelated, even though the presupposition itself remains constant.

4. Experiments

We evaluated four language models to determine whether they process presuppositions in conditionals according to human judgments or theory-based predictions. Our evaluation combines classification accuracy with explainability analyses across controlled varied datasets. Human labels reflect empirical judgments from the original examples, while theory-based labels follow the conditional projection patterns predicted by formal semantic theories discussed in Section 1. Below, we describe the models, evaluation metrics, and experimental procedures.

4.1. Models and Setup

Our experiments use RoBERTa-large-MNLI, DeBERTa-large-MNLI, Llama-3.2-1B, and Gemma-3-1B. RoBERTa and DeBERTa were selected as they are NLI-specific models, pre-trained on MultiNLI (Williams et al., 2018). Llama and Gemma were included as they are recent, open-weight

Subset	Premise	Hypothesis	Human	Theory
1	DEP: If John is a scuba diver, he'll bring his wetsuit. IND: If Lisa finishes meeting early, she'll never drive a sports car again.	John has a wetsuit. Lisa has driven a sports car before.	E	N
2	Conjunction: If Sarah attends the conference and her friend meets her at the hotel, it's okay. Disjunction: Either Sarah doesn't attend conference, or she and her friend do. Belief: Sarah believes her friend is upstairs.	Sarah has a friend.	E	N
3	Related: If Nadia attends conference, she'll never watch Star Wars again. Somewhat related: If Nadia attends conference, she'll never watch a George Lucas movie again. Unrelated: If Nadia attends conference, she'll never watch Titanic again.	Nadia has watched Star Wars.	E/N*	N
4	Related: If Ali attends conference, his brother will meet her at university. Unrelated: If Ali attends conference, his brother will get rid of his old car.	Ali has a brother.	E	N

Table 2: Examples from the four subsets with corresponding human and theory-based labels. *In Subset 3, the human label is E for the related variant where the trigger matches the hypothesis, and N for somewhat related and unrelated variants where the modified trigger no longer semantically supports the hypothesis.

large-language models. All models were fine-tuned on the CONFER training set for one epoch with a learning rate of $5e-5$.

4.2. Evaluation Metrics

We use both classification accuracy and explainability metrics to assess whether model predictions align with human or theory-based labels and to determine which linguistic features drive those predictions.

Classification Accuracy is calculated against both human labels and theory-based labels.

Integrated Gradients (IG) (Sundararajan et al., 2017) measures token-level attribution by integrating the gradients of the model output along a path from a baseline input x' to the actual input x :

$$IG_i(x) = (x_i - x'_i) \int_0^1 \frac{\partial F(x' + \alpha(x - x'))}{\partial x_i} d\alpha \quad (1)$$

where x is the input, x' the baseline, x_i a token, $\alpha \in [0, 1]$ the interpolation factor, and $\frac{\partial F}{\partial x_i}$ the partial derivative of the model output with respect to token i .

IG applies the same backpropagation mechanism that the model uses during optimization to compute gradients of the output with respect to the input. Like perturbation-based methods such as LIME (Ribeiro et al., 2016) and SHAP (Lundberg and Lee, 2017), it identifies important input features, but does so more cost-effectively by using gradient information rather than iterative input modifications (Bastings and Filippova, 2020).

Using this method, we can address an important question: do models focus on the same linguistic elements that humans consider crucial for presuppositional reasoning? In presuppositional entailment,

trigger phrases provide the critical link between conditionals and their presuppositions. For example, in *If Randolph is a carpenter, he'll use his beading tools for designing*, the phrase *his beading tools* triggers the inference that Randolph has beading tools. Without this trigger, the entailment would not arise. Figure 2 illustrates how IG highlights such tokens.

From the IG scores, we derive another additional metric called the trigger IG ratio, which quantifies the proportional influence assigned to the trigger words relative to the average influence of all tokens.

Context-Trigger Attention Metrics complement IG by showing how models encode relationships between phrases rather than individual token salience. While IG shows which tokens matter for predictions, attention patterns indicate whether models connect presupposition triggers with their contextual dependencies, which is essential for pragmatic reasoning in conditionals.

For Subset 4, we compute two metrics:

- **K1→K2 attention:** measures attention flow from the noun phrase preceding the presupposition trigger in the hypothesis (K1) to the contextual phrase (K2), normalized by overall attention. This assesses whether models establish stronger connections between semantically related elements.
- **K2→special token attention:** measures attention from K2 to delimiter tokens ([SEP], <bos>, [CLS]), normalized by average special token attention. This indicates whether models treat context as structurally significant for sentence representation.

T-test for Significance assesses whether modifications significantly affect trigger IG ratios. We

If Theo takes the train to the airport, his wife will meet him at the gate.
Theo has a wife

Figure 2: IG visualization showing token-level attribution. Darker shading indicates higher IG values, with the presupposition trigger receiving the strongest attribution.

```
Consider the truth value of the premise.  
If the premise is true, does it necessarily  
mean that the hypothesis must also be true?  
  
Output E if the hypothesis must always be true.  
Output C if the hypothesis must always be false.  
Output N if the hypothesis may be either  
true or false.  
  
Do not output anything other than E, C, or N.  
  
Premise: {premise}  
Hypothesis: {hypothesis}  
Output:
```

Figure 3: Three-way classification prompt used for zero-shot evaluation of Gemma and LLaMA models.

perform paired t-tests (reported with p-values) comparing original and modified examples ($\alpha = 0.05$), indicating which structural and semantic variations reliably affect model attention patterns. FDR multiple-comparison correction was applied.

4.3. Zero-Shot Evaluation

To assess whether models can handle presuppositional reasoning based on their pretraining without task-specific fine-tuning, we evaluated all four models on subsets 1-4 in a zero-shot setting. RoBERTa and DeBERTa, pre-trained on MultiNLI, were tested directly as a three-way classification task (Entailment, Neutral, Contradiction). Gemma and LLaMA were evaluated through zero-shot prompting using the instruction shown in Figure 3. We computed classification accuracy against both human and theory-based labels to establish baseline performance before fine-tuning.

4.4. Structural Variations

To examine whether structural variations affect model sensitivity to presupposition triggers, we fine-tuned all four models on CONFER and evaluated them on Subset 2, which contains three structural modifications explained in 3.2. We computed classification accuracy alongside the trigger IG ratio to assess whether these structural changes shift the importance models place on presuppositional content.

4.5. Lexical-Semantic Substitutions

Experiments on Subset 3 evaluate whether models process presupposition triggers based on semantic content or surface-level position. As described in Section 3.3, we created three relatedness levels by substituting trigger phrases with semantically related, somewhat related, or unrelated alternatives, updating gold labels from Entailment to Neutral when the semantic relationship no longer holds. We measured classification accuracy and trigger IG ratio scores to assess whether models appropriately adjust their predictions when trigger semantics change, or continue to rely on positional heuristics.

4.6. Contextual Modifications

To assess whether models are sensitive to the logical relationship between conditional components, we evaluated fine-tuned models on Subset 4, which manipulates premise-level context while keeping the presupposition trigger constant. As described in Section 3.4, we created variants where the antecedent and consequent events are either logically related or unrelated. We measured classification accuracy alongside $K1 \rightarrow K2$ attention and $K2 \rightarrow$ special token attention to determine whether models maintain appropriate focus on presuppositional triggers regardless of contextual distractors, or whether unrelated context phrases interfere with presupposition processing.

5. Results and Analysis

5.1. Models Align with Human Judgments and Not Theories

In zero-shot evaluation, we assessed whether pre-trained models exhibit human-like or theory-aligned presuppositional reasoning. As indicated in Figure 5, RoBERTa and DeBERTa achieved perfect accuracy against human labels on original examples (Subset 1). In contrast, they showed 0% accuracy against theory-based labels. The results indicate complete alignment with human unconditional interpretations rather than the conditional presuppositions predicted by formal semantics. In this experiment, LLaMA and Gemma underperformed compared to RoBERTa and DeBERTa. LLaMA achieved 74-86% accuracy against human labels,

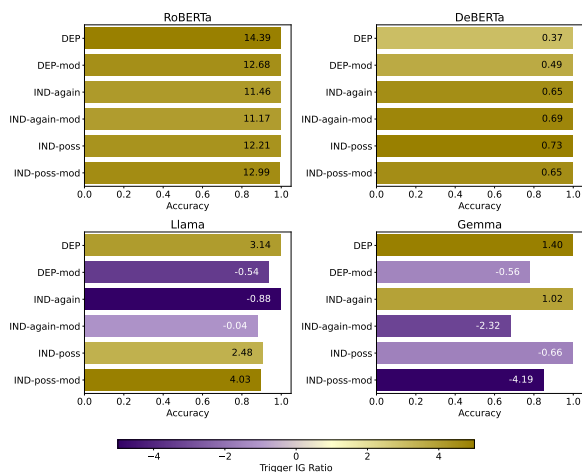


Figure 4: Accuracy across sentence types in Subset 2. The horizontal axis shows accuracy percentages, and the values inside the bars indicate the corresponding trigger IG ratios.

while Gemma struggled across all sentence types in this subset.

On modified examples (Subsets 2-4), models maintained alignment with human judgments for Subsets 2 (structural variations) and 4 (contextual modifications), with accuracy remaining above 90% for RoBERTa and DeBERTa. However, performance dropped dramatically on Subset 3 (semantic substitutions), where models correctly identified only 24-52% of examples as N despite the broken semantic relationship between modified triggers and hypotheses. This shows that models rely on structural and positional cues rather than semantic content when processing presupposition triggers.

5.2. Structural Modifications Maintain Trigger Focus Despite Complexity

Fine-tuned models maintained high classification accuracy on Subset 2, with RoBERTa and DeBERTa achieving 99-100% accuracy on both original and modified examples. However, models showed different attention patterns. RoBERTa indicated the highest trigger IG ratios across both original and structurally embedded examples, indicating strong focus on presupposition triggers regardless of syntactic complexity (Figure 4). Statistical tests showed no significant differences in trigger IG ratios between original and modified examples ($p > 0.05$), suggesting RoBERTa's trigger sensitivity is robust to structural variations.

DeBERTa showed lower trigger IG ratios (< 1) across all examples, indicating below-average attention to triggers, though accuracy remained nearly perfect. LLaMA and Gemma showed greater variability in trigger IG, including occasional negative values, with modest accuracy drops.

5.3. Models Rely on Trigger Position Over Semantic Content

Models achieved high accuracy on Subset 3's unmodified original examples but failed dramatically when trigger phrases were semantically modified. Despite human labels changing to N, models continued predicting E, showing reliance on structural position over semantic content.

Performance varied by trigger type and model. IND-again examples proved most difficult, with all models achieving below 4% accuracy, nearly complete misclassification. On DEP-poss examples, models performed marginally better with RoBERTa reaching 22% accuracy while others remained below 15%. IND-poss examples were relatively easier: RoBERTa and Gemma achieved above 35% accuracy, and LLaMA surprisingly exceeded 80%. However, DeBERTa misclassified all IND-poss examples.

Trigger IG analysis confirmed models' reliance on structural position (Figure 6). Across all forms, significant differences did not emerge between the way models treat trigger phrases in the original examples compared to the modified examples ($p > 0.05$), despite semantic changes. RoBERTa maintained high IG ratios regardless of semantic fit, while DeBERTa showed consistently low values, with ratios stable between 0.8-0.85 even in IND-again examples. LLaMA and Gemma exhibited variability, with LLaMA occasionally increasing IG in unrelated contexts and Gemma showing negative values. Moreover, LLaMA achieving 80% accuracy on IND-poss examples, does see a slight negative IG shift, correctly indicating that the trigger phrase no longer carries importance in the neutral relationships. The overall stable attribution despite semantic manipulation confirms that models rely on trigger position rather than meaning.

5.4. Contextual Modifications Expose Training Set Overfitting

Models maintained perfect accuracy ($\sim 100\%$) on modified DEP and IND-poss examples in Subset 4. They correctly predicted E regardless of whether contextual phrases (K2) were semantically related or unrelated to the noun phrase in the hypothesis (K1). However, accuracy decreased significantly on modified IND-again examples, with LLaMA achieving only 47% and the other three models below 15% (Figure 7). Models incorrectly predicted N instead of E despite only irrelevant context words changing. Since zero-shot models achieved $> 80\%$ accuracy on these same examples (except Gemma), this failure appears to result from CONFER fine-tuning. CONFER includes examples where semantically related antecedents pair with *again* triggers but carry N labels due to superset relationships in the




Figure 5: Models' accuracies in zero-shot evaluation across four subsets, using human and theory-based labels.




Figure 6: Accuracy and trigger IG ratios across sentence types in Subset 3

antecedent. Models likely learned the spurious pattern "related antecedent + again = N" from these examples, causing confusion with the structurally similar IND-again modifications.

Attention analysis indicates distinct patterns by trigger type. $K1 \rightarrow K2$ attention showed significant variation across all trigger types. For DEP and IND-poss examples, these attention shifts corresponded with models' accuracy patterns. In IND-again examples, models placed significantly less attention from the trigger noun phrase to context in modified variants compared to original ones, showing the substantial accuracy drops observed in this condition.

Models also significantly increased $K2 \rightarrow$ special token attention in modified IND-again examples compared to originals ($p < 0.08$ for RoBERTa, $p < 0.001$ for others), indicating these contextual modifications affected sentence-level representations.




Figure 7: Subset 4 accuracy by sentence type and model, with bar colours indicating $K2 \rightarrow$ special token attention on the corresponding examples.

While RoBERTa and Gemma showed some attention shifts in other subsets, only IND-again modifications produced consistent, significant changes across all four models, the same subset where accuracy decreased. This correlation between shifted attention patterns and misclassification shows overfitting to training patterns where "related context + again" was associated with N labels, leading models to apply this association inappropriately. Figure 8 illustrates the average attention between tokens in Gemma.

6. Discussion and Conclusion

We introduced a diagnostic resource for evaluating how language models process presuppositional reasoning in conditionals, addressing the proviso

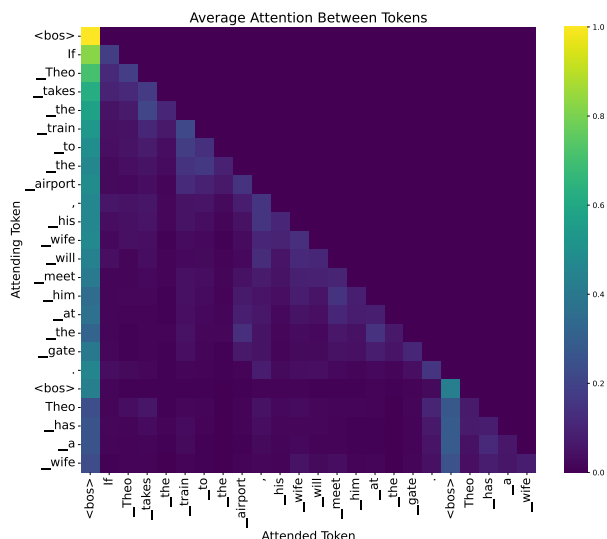


Figure 8: Average attention between tokens across all layers and heads in Gemma, for the sentence *If Theo’s train is delayed, his wife will wait for him at the station*, and its hypothesis *Theo has a wife*. Brighter cells indicate higher attention values. First column corresponds to the special token attention.

problem. We created a dataset of 8,500 examples with controlled linguistic variations and an evaluation framework combining classification accuracy with explainability methods. This framework, facilitates assessment of whether models process presuppositions through semantic understanding or shallow heuristics, with comparisons to human reasoning patterns.

Our evaluation shows that while models align with human judgments, they achieve this through pattern matching rather than semantic reasoning. Diagnostic tests revealed three main limitations. First, models dramatically fail when trigger phrases are replaced with semantically unrelated alternatives. Second, accuracy alone can be misleading for assessing model reasoning. On Subset 2, RoBERTa achieves 99% accuracy by focusing more on trigger words (IG scores 11-17), while DeBERTa achieves identical accuracy while putting minimal importance on the trigger (IG < 1). However, Subset 3 then reveals the extent to which DeBERTa is fallible under the modifications, achieving an accuracy of 0%, while RoBERTa is more robust, still capable of achieving an accuracy between 20% and 50% on Subset 3. Third, fine-tuning can introduce misleading correlations that degrade performance on structurally similar (but logically different) examples.

This work demonstrates that evaluating pragmatic reasoning in language models requires diagnostic resources paired with multi-method assessment frameworks. While accuracy metrics suggest strong performance, our explainability analysis

shows models rely on structural patterns rather than semantic understanding. Future work can extend this diagnostic approach to other pragmatic phenomena such as scalar and conventional implicature. Moreover, exploring task formats beyond NLI would assess whether findings generalize across evaluation paradigms. Additionally, parallel psycholinguistic experiments could test longstanding semantic theories empirically by directly comparing human processing patterns with model behavior.

7. Limitations

This study focuses on a specific type of conditional sentences structured in the form *If A, B_p*, where the presupposition trigger appears in the consequent. The conditionals include only two trigger types: possessive pronouns and the adverbial trigger *again*. These triggers were selected because they occur frequently in the source CONFER dataset (Azin et al., 2025) and yield presuppositions that can be directly restated as standalone hypotheses, making them suitable for the NLI formulation used in our experiments. This design choice means the dataset does not capture the full range of presuppositional phenomena in natural language, including other trigger types (e.g., factive verbs, aspectual verbs, clefts), different conditional constructions, or interactions with broader discourse context.

Our evaluation is limited to four language models, and findings may not generalize to other model architectures, sizes, or training paradigms. The NLI task format, while appropriate for testing inferential relationships, represents only one way of probing presuppositional reasoning and may not fully reflect how models process presuppositions in generation or other downstream tasks.

8. Ethical Considerations

All experimental materials in this study consist of everyday scenarios with no sensitive content, stereotypes, or references to real individuals. The sentences were designed only to test linguistic phenomena. As this is a diagnostic evaluation study rather than a deployment of models for real-world applications, the research poses minimal ethical risks. The dataset and code will be made publicly available upon publication for transparency and reproducibility.

9. Bibliographical References

Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. [A diagnostic study of explainability techniques for text](#)

- classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3256–3274, Online. Association for Computational Linguistics.
- Katherine Atwell, Mandy Simons, and Malihe Alikhani. 2025. [Measuring bias and agreement in large language model presupposition judgments](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 2096–2107, Vienna, Austria. Association for Computational Linguistics.
- Tara Azin, Daniel Dumitrescu, Diana Inkpen, and Raj Singh. 2025. Let’s confer: A dataset for evaluating natural language inference models on conditional inference and presupposition. In *Proceedings of the Canadian Conference on Artificial Intelligence (CAIAC)*. Retrieved from <https://caiac.pubpub.org/pub/keh8ij01>.
- Jasmijn Bastings and Katja Filippova. 2020. [The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?](#) In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, Online. Association for Computational Linguistics.
- David I. Beaver. 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, CA.
- David Ian Beaver. 1997. Presupposition. In Johan Van Benthem and Alice Ter Meulen, editors, *Handbook of Logic and Language*, chapter 17, pages 939–1008. North-Holland.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen tau Yih, and Yejin Choi. 2020. Abductive common-sense reasoning. In *Proceedings of the 2020 International Conference on Learning Representations (ICLR)*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642. Association for Computational Linguistics.
- Andre Cianflone, Yulan Feng, Jad Kabbara, and Jackie Chi Kit Cheung. 2018. Let’s do it “again”: A first computational approach to detecting adverbial presupposition triggers. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2747–2755, Melbourne, Australia. Association for Computational Linguistics.
- Robin Cooper, Richard Crouch, Jan Van Eijck, Chris Fox, Josef Van Genabith, Jan Jaspers, Hans Kamp, Manfred Pinkal, Massimo Poesio, and Stephen Pulman. 1994. *FraCaS: A Framework for Computational Semantics*. Deliverable D6.
- Marie-Catherine de Marneffe, Mandy Simons, and Judith Tonhauser. 2019. The commitmentbank: Investigating projection in naturally occurring discourse. In *Proceedings of Sinn und Bedeutung*, volume 23, pages 107–124.
- Judith Degen. 2015. Investigating the distribution of some (but not all) implicatures using corpora and web-based methods. *Semantics and Pragmatics*, 8:11–1.
- Bart Geurts. 1996. Local satisfaction guaranteed: A presupposition theory and its problems. *Linguistics and Philosophy*, 19:259–294.
- Bart Geurts. 1999. *Presuppositions and Pronouns*. Elsevier, Amsterdam.
- Irene Heim. 1983. On the projection problem for presuppositions. In Paul Portner and Barbara H. Partee, editors, *Formal Semantics: The Essential Readings*, pages 249–260. Blackwell.
- Irene Heim. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3):183–221.
- Wesley H. Holliday, Matthew Mandelkern, and Cedegao E. Zhang. 2024. [Conditional and modal reasoning in large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3800–3821, Miami, Florida, USA. Association for Computational Linguistics.
- Paloma Jeretič, Alicia Parrish, Sheng-Fu Wang, Jason Phang, Anhad Mohananey, Phu Mon Htut, and Samuel R. Bowman. 2019. Investigating BERT’s knowledge of language: Five analysis methods with NPIs. 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 2870–2880.
- Paloma Jeretic, Alex Warstadt, Suvrat Bhooshan, and Adina Williams. 2020. [Are natural language inference models IMPPRESSive? Learning IMPLicature and PRESupposition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8690–8705, Online. Association for Computational Linguistics.

- Nanjiang Jiang and Marie-Catherine de Marneffe. 2019. Do you know that florence is packed with visitors? evaluating state-of-the-art models of speaker commitment. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4208–4213, Florence, Italy. Association for Computational Linguistics.
- Lauri Karttunen. 1973. Presuppositions of compound sentences. *Linguistic Inquiry*, 4(2):169–193.
- Lauri Karttunen and Stanley Peters. 1979. Conventional implicatures in Montague grammar. In Choon-Kyu Oh and David A. Dinneen, editors, *Syntax and Semantics 11: Presupposition*, pages 1–56. Academic Press, New York, USA.
- David Lewis. 1979. [Scorekeeping in a language game](#). *Journal of Philosophical Logic*, 8(1):339–359.
- Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 4768–4777. Curran Associates Inc.
- Matthew Mandelkern. 2016. Dissatisfaction theory. In *Proceedings of the 26th Semantics and Linguistic Theory (SALT)*, pages 391–416.
- Clemens Mayr and Jacopo Romoli. 2016. Satisfied or exhausted: An ambiguity account of the proviso problem. In *Proceedings of Semantics and Linguistic Theory (SALT)*, pages 892–912.
- George A. Miller. 1995. Wordnet: A lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Alyssa Parrish, Sebastian Schuster, Alex Warstadt, et al. 2021. [Nope: A corpus of naturally-occurring presuppositions in english](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 349–366.
- Christopher Potts. 2015. [Presupposition and implicature](#). In *The Handbook of Contemporary Semantic Theory*, pages 168–202. John Wiley & Sons, Ltd.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ["why should i trust you?": Explaining the predictions of any classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Alexis Ross and Ellie Pavlick. 2019. How well do nli models capture verb veridicality? 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 2230–2240.
- Sebastian Schuster, Yuxing Chen, and Judith Degen. 2020. Harnessing the richness of the linguistic signal in predicting pragmatic inferences. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Raj Singh. 2007. [Formal alternatives as a solution to the proviso problem](#). In *Semantics and Linguistic Theory (SALT) 17*.
- Raj Singh. 2020. Matrix and embedded presuppositions. In *The Wiley Blackwell Companion to Semantics*, pages 1–42. Wiley-Blackwell.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, pages 4444–4451. AAAI Press.
- Robert Stalnaker. 1973. [Presuppositions](#). *Journal of Philosophical Logic*, 2(4):447–457.
- Robert Stalnaker. 1998. On the representation of context. *Journal of Logic, Language and Information*, 7(1):3–19.
- Peter F. Strawson. 1950. On referring. *Mind*, 59(235):320–344.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 3319–3328. PMLR.
- Kai von Fintel. 2008. [What is presupposition accommodation, again?](#) *Philosophical Perspectives*, 22(1):137–170.
- Aaron Steven White and Kyle Rawlins. 2018. The role of veridicality and factivity in clause selection. In *Proceedings of the 48th Annual Meeting of the North East Linguistic Society*, Amherst, MA, USA. GLSA Publications.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.

Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104. Association for Computational Linguistics.

10. Language Resource References

Tara Azin and Daniel Dumitrescu and Diana Inkpen and Raj Singh. 2025. *CONFER: A Dataset for Evaluating Natural Language Inference Models on Conditional Inference and Presupposition*. Canadian Artificial Intelligence Association (CAIAC). PID <https://github.com/Conditional-NLI/CONFER>.

George A. Miller. 1995. *WordNet: A Lexical Database for English*. Princeton University. PID <https://wordnet.princeton.edu/>.

Robyn Speer and Joshua Chin and Catherine Havasi. 2017. *ConceptNet 5.5: An Open Multilingual Graph of General Knowledge*. Commonsense Computing Initiative. PID <https://github.com/commonsense/conceptnet5>.

Adina Williams and Nikita Nangia and Samuel Bowman. 2018. *MultiNLI: A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference*. New York University. PID https://huggingface.co/datasets/nyu-ml/multi_nli.