

Are BabyLMs Deaf to Gricean Maxims?

A Pragmatic Evaluation of Sample-efficient Language Models

Raha Askari^{1,2}, Sina Zarrieß², Özge Alacam², Judith Sieker²

¹Department of Humanities, University of Turin

²Computational Linguistics, Department of Linguistics, Bielefeld University

Abstract

Implicit meanings are integral to human communication, making it essential for language models to be capable of identifying and interpreting them. Grice (1975) proposed a set of conversational maxims that guide cooperative dialogue, noting that speakers may deliberately violate these principles to express meanings beyond literal words, and that listeners, in turn, recognize such violations to draw pragmatic inferences. Building on Surian et al. (1996)’s study of children’s sensitivity to violations of Gricean maxims, we introduce a novel benchmark to test whether language models pretrained on <10M and <100M tokens can distinguish maxim-adhering from maxim-violating utterances. We compare these BabyLMs across five maxims and situate their performance relative to children and a Large Language Model (LLM) pretrained on 3T tokens. We find that overall, models trained on <100M tokens outperform those trained on <10M, yet fall short of child-level and LLM competence. Our results suggest that modest data increases improve some aspects of pragmatic behavior, leading to finer-grained differentiation between pragmatic dimensions.

Our benchmark extends the BabyLM evaluation suite to pragmatic aspects of language and is publicly available.¹

1 Introduction

Consider the following exchange: Sarah asks her friend “What did you eat for lunch?”, upon which her friend might reply “I had something edible” or “I had chicken soup with an extra small silver spoon”. While both responses are true and perfectly grammatical, the first one fails to provide the amount of information Sarah’s question calls for, and the second one contains excessive, unasked details. Most listeners would expect an answer that

is specific but not unnecessarily detailed, such as “I had chicken soup”.

In everyday conversation, such under- or over-informative replies stand out as odd because they do not provide the adequate amount of detail the question asks for. The philosopher Grice (1975) explained such phenomena through his *Cooperative principle*, which holds that speakers are generally aware of what is conversationally suitable or unsuitable. He proposed a set of conversational maxims, one of which, the *maxim of Quantity*, requires the speaker to be as informative as necessary.

The ability to notice and interpret such deviations from conversational norms is a key aspect of pragmatic competence, and essential for successful communication. In the evaluation of Large Language Models (LLMs), however, while models are now routinely tested on a wide range of syntactic tasks (e.g., Marvin and Linzen (2018); Hu et al. (2020); Finlayson et al. (2021); Lampinen (2024); Kryvosheieva and Levy (2025)), far fewer studies target their ability to reason pragmatically (Ettinger, 2020; Fried et al., 2023; Ma et al., 2025a; Sieker et al., 2025; Lachenmaier et al., 2025). This gap is especially pronounced for resource-limited models such as those developed for the BabyLM Challenge (Warstadt et al., 2023; Choshen et al., 2024). One possible reason is that, unlike syntax, pragmatics does not easily lend itself to large-scale minimal-pair test creation. Controlled operationalizations of pragmatic phenomena, such as those in Sieker et al. (2023) and Sieker and Zarrieß (2023), remain rare and resource-intensive, highlighting the challenge of designing systematic evaluation materials for this domain.

In psycholinguistics, however, several diagnostic tasks for pragmatic understanding already exist (e.g., Doran et al. (2012); Degen and Tanenhaus (2014); Romoli and Schwarz (2015); Tieu et al. (2015)). One such task is the Conversational Violations Test (CVT), designed to investigate children’s

¹https://huggingface.co/datasets/rahaaskari/gricean_baby

pragmatic abilities based on Gricean maxims, introduced in Surian et al. (1996)’s study "Are Children with Autism Deaf to Gricean Maxims?". In the CVT, children are presented with short dialogues where one answer follows one of Gricean conversational maxims and another violates it. Children are asked to identify the maxim-violating response. This controlled forced-choice format makes the CVT particularly attractive for LM evaluation: the correct choice depends on recognizing conversational norms rather than relying solely on factual knowledge or grammar, and the task fits well to established evaluation methods that compare model-assigned probabilities for a predefined set of candidate responses, as in grammatical acceptability (Warstadt et al., 2020), abductive commonsense reasoning (Zhao et al., 2023), or semantic relations of compound nouns (Rambelli et al., 2024).

In this paper, we adapt the CVT into a benchmark for evaluating BabyLMs’ sensitivity to Gricean maxims. Starting from the original 25 conversational items from CVT, we augment the dataset automatically to over 2,250 items and refine them through human annotation. We evaluate a range of BabyLM baseline models (4 trained on <10M and 4 trained on <100M words), compare their performance to that of children from Surian et al. (1996), and situate their results alongside an LLM trained on more extensive data. In total, our experiment produces 20,250 data points (across 8 BabyLMs and 1 LLM). Among the evaluated models, BabyLMs trained on <100M tokens outperformed those trained on <10M, yet both groups fell short of achieving child-level pragmatic accuracy despite their developmental motivation. On average, BabyLMs performed best when judging truthfulness, but struggled most with assessing the appropriate level of informativeness. The LLM generally surpassed the BabyLMs and, in some cases, even outperformed children, but still failed to match children’s overall competence. Thus, despite vastly larger training data, notable gaps persist between model and child performance across several maxims.

The contributions of this study are threefold: (1) a novel, linguistically controlled benchmark for testing pragmatic competence in data-limited LMs, grounded in established psycholinguistic methodology; (2) an empirical analysis of BabyLMs’ performance across different Gricean maxims; and (3) a comparison of LM and child performance that situates model behavior within a developmental

trajectory.

2 Background

Effective communication relies on more than just producing grammatical sentences. Much of what speakers communicate is conveyed implicitly, relying on the listener to infer meanings that go beyond the literal words. To do so successfully, speakers must choose utterances that are appropriate to the conversational context, and listeners must interpret them in light of shared assumptions, intentions and social norms. Even a perfectly well-formed sentence can be unhelpful, misleading or socially awkward if it ignores these unspoken rules. The study of how meaning is shaped by such contextual factors is the domain of pragmatics, which rests on the central idea that conversation is a cooperative activity: participants work together to exchange information efficiently and meaningfully. Grice (1975) formalized this intuition in his *Cooperative Principle*, according to which interlocutors are generally aware of what is conversationally suitable or unsuitable at each stage of a dialogue. He categorized this principle into four maxims, and additionally discussed Politeness as what he termed an "off-the-list" maxim:

- Quantity (*Make your contribution as informative as required and do not make your contribution more informative than is required*)
- Quality (*Do not say what you believe to be false and do not say that for which you lack adequate evidence*)
- Relation (*Be relevant*)
- Manner (*Be perspicuous, i.e., avoid obscurity, avoid ambiguity, be brief and be orderly*).
- Politeness (*Be polite*)

While these maxims are typically adhered to, speakers may sometimes blatantly violate them by saying one thing but implying another, producing what is known as an *implicature*. For example, when two colleagues are talking during a lunch break, one might ask "Did you talk to the boss about the promotion?", and the other might reply, "I really like this food." This response violates the maxim of Relation and prompts the listener to search for the intended meaning, assuming the other person remains cooperative and aware of the maxims. In this case, for example, the interlocutor

is likely to infer that their colleague wishes to avoid the topic and has not spoken to the boss. Over the past decades, the Gricean maxims have become a cornerstone of pragmatic theory, shaping how researchers analyze and explain the ways people interpret and produce language in context.

Developmental Studies. Several developmental psycholinguistic studies have examined the age at which the sensitivity to such conversational violations emerges in humans (to name a few; Ackerman, 1981; Conti and Camras, 1984; Surian et al., 1996; Surian et al., 2010; Okanda et al., 2015 and Panzeri and Foppolo, 2021). In this direction, Surian et al. (1996)’s study introduced the Conversational Violations Test (CVT) to compare the pragmatic abilities of children with autism and specific language impairments to those of neurotypical children by incorporating Grice’s framework. The maxims addressed in their study were Quantity (divided into two maxims; I: *Be informative* and II: *Avoid redundant information*), Quality, Relation and Politeness. The CVT is a set of 25 short conversational items and contains 5 conversations for each maxim. In their experiment, 8 neurotypical children (mean age 6-7) were presented with tape-recorded conversations featuring three puppets. In each scenario, one puppet would ask a question, and the other two would respond, only that one of them would provide an answer that violated a conversational maxim. The children were then asked to identify the puppet that gave the *silly or funny* answer, i.e., the one that violated the maxim. See Table 1 for examples for each maxim.

BabyLMs. The BabyLM Challenge aims to model human language development in order to optimize language model pretraining under strict data limitations. Tracks for submissions of text-only models include the Strict-small track (trained on <10M tokens) and the Strict track (trained on <100M tokens). As a starting point for evaluation, the organizers release a group of baseline models accompanied with pretraining corpora, along with an evaluation pipeline including benchmarks such as BLiMP (Warstadt et al., 2020) or GLUE (Wang et al., 2019)². While these benchmarks provide broad coverage of linguistic competence, they do not directly and comprehensively assess pragmatic abilities. Given the developmental motivation be-

²Find a complete overview of the BabyLM evaluation pipeline at <https://github.com/babylm/evaluation-pipeline-2025>.

Maxim	Example
Quantity I <i>Be informative</i>	Q: <i>How do you prefer your tea?</i> Follower: <i>With milk.</i> Violator: <i>In a cup.</i>
Quantity II <i>Avoid redundant information</i>	Q: <i>Who is your best friend?</i> Follower: <i>My best friend is John. He goes to my school.</i> Violator: <i>My best friend is Peter. He wears clothes.</i>
Quality <i>Be truthful</i>	Q: <i>Where do you live?</i> Follower: <i>I live in London.</i> Violator: <i>I live on the moon.</i>
Relation <i>Be relevant</i>	Q: <i>What games do you know?</i> Follower: <i>I know how to play football.</i> Violator: <i>I know your name.</i>
Politeness <i>Be polite</i>	Q: <i>Do you like my dress?</i> Follower: <i>It’s pretty.</i> Violator: <i>I hate it.</i>

Table 1: Example items for different conversational maxims from Surian et al. (1996)’s CVT. The *Follower* adheres to the maxim, while the *Violator* does not.

hind BabyLMs, we argue that it is equally important to examine whether such models exhibit the pragmatic reasoning observed in humans.

Pragmatic Evaluation in LMs. As pragmatic knowledge is essential for successful communication, recent studies have explored whether LLMs exhibit pragmatic reasoning. Some studies report that LLMs can perform competitively with humans on specific tasks such as metaphor comprehension (Hu et al., 2023; Sanchez-Bayona and Agerri, 2025), but many find that they still struggle with a wide range of phenomena, including sarcasm and jokes (Hu et al., 2023; Jentsch and Kersting, 2023), theory of mind (Shapira et al., 2023; Trott et al., 2023; Gandhi et al., 2024), implicit causality (Sieker et al., 2023; Kankowski et al., 2025), context-dependent reference resolution (Junker et al., 2025; Ma et al., 2025b), or inferences like presuppositions (Kabbara and Cheung, 2022; Sieker and Zarri , 2023; Tsvilodub et al., 2024; Sieker et al., 2025; Lachenmaier et al., 2025).

When it comes to the Gricean maxims, Hu et al. (2023), for example, evaluated LLMs’ ability to understand intended meanings by prompting models with short English stories and asking what a char-

acter wanted to convey by flouting a maxim, given a set of possible answers. They found that the models would generally assign higher probabilities to literal meanings over the speaker’s intended meaning. Similarly, LLMs demonstrated bias towards literal meanings during a pragmatic evaluation for Korean language by Park et al. (2024). In their experiment, models performed poorly on the maxim of Relation and well on the maxim of Quality when selecting the pragmatic meanings from given options, but showed reversed patterns for open-ended questions about the speaker’s intent. Yue et al. (2024), on the other hand, evaluated models’ ability to infer implicated meanings in multi-turn Chinese dialogues and found no significant variation in model performances across maxims. Moreover, most models failed to generate correct interpretations for implicatures despite being able to identify them in a multiple-choice setting. Other examples of pragmatic evaluations of LLMs by incorporating implicatures include (but are not limited to) Zheng et al. (2021), who presented the GRICE dataset for assessing the pragmatic reasoning of LLMs while taking into account other aspects of modern dialogue modeling like coreference; Cho and Kim (2024), who compared cosine similarities of literal meanings of scalar implicatures with their pragmatic meanings; and Kurch et al. (2024), who tested whether LLMs can derive atypicality inferences that are triggered through information redundancy.

Building on this line of work, we extend prior studies on LM pragmatic competence and Gricean maxims by introducing a child-directed, maxim-balanced benchmark that enables direct comparison between model and child performance. Inspired by Surian et al. (1996)’s CVT, we compile a dataset of 2,250 conversational items in a controlled forced-choice format. Our benchmark is particularly well-suited to the BabyLM Challenge because its simple, child-appropriate language and controlled design offer a fine-grained, diagnostic test of pragmatic abilities, while minimizing reliance on large-scale training data or extensive world knowledge.

3 Approach

In order for pragmatic interpretations (those that go beyond literal ones) to arise, a listener must know the rules of conversation, recognize when they are violated, and discern when an utterance may be literally unfitting (uncooperative) yet prag-

matically acceptable (cooperative). Building on the framework of Gricean maxims, Surian et al. (1996) examined the pragmatic competence of children by testing whether they could identify an uncooperative (i.e., maxim-violating) answer among a pair of responses to a given question. We adopt this same forced-choice paradigm to evaluate the pragmatic sensitivity of language models.

Data. We base our evaluation on the CVT and extend this resource into a large-scale benchmark by generating additional CVT-style items with GPT-4 (OpenAI, 2024) and manually curating the outputs to maintain child-level vocabulary³, grammaticality, naturalness and adherence to the targeted maxim. The final dataset contains 2,250 dialogues, balanced across five maxims: Quantity I, Quantity II, Quality, Relation, and Politeness. Full details of the augmentation process and quality control criteria are provided in Appendix A.1. Also, see Appendix A.2 for examples of experimental items of our dataset.

Models. We use the following baseline BabyLMs pretrained on BabyLM corpora that were released in two tracks (Strict for models trained on at most 100M tokens and Strict-small for models trained on at most 10M tokens)⁴: two auto-regressive LMs, namely GPT-2 (Radford et al., 2019) and Baby Llama (Timiryasov and Tastet, 2023) and two masked LMs, namely LTG-BERT (Samuel et al., 2023) and Roberta (Liu et al., 2019). Finally, to assess the effect of more training data on the pragmatic performance of language models and to enable a comparison with an LLM, we evaluate the decoder-only OLMo-1B (Groeneveld et al., 2024) as a representative of fully open LLMs that has been trained on 3T tokens.

Evaluation. Using our curated dataset, we evaluate pragmatic sensitivity of language models in an unsupervised setting. Specifically, we measure a model’s sentence acceptability for the two candidate answers to a question: one that follows a Gricean maxim (follower) and one that violates it (violator). For incremental models, we compute the conditional log-probability of the answer given the question, while in the case of

³This is derived from the fact that the pretraining data for baseline BabyLMs mostly consists of input received by children.

⁴Baseline models for previous years and this year’s submission are available at <https://huggingface.co/babylm> and <https://huggingface.co/BabyLM-community>.

masked language models, we use the improved pseudo-log-likelihood proposed by Kauf and Ivanova (2023). In both cases, the probability of the answer is calculated as the sum of the log-probabilities of its tokens, normalized by its length. For each item, we assess whether the model assigns a higher probability to the maxim-following answer:

$$\mathbb{1}[P(\text{Answer}_{\text{Follower}} | \text{Question}) > P(\text{Answer}_{\text{Violator}} | \text{Question})]$$

Model accuracy for each maxim is defined as the proportion of items for which the model assigns a higher probability to the maxim-follower response. We obtain models’ scores through Minicons (Misra, 2022)⁵, which is an open-source library for extracting sentence acceptability measures in language models.

4 Results

In Table 2, we report accuracy per maxim for the BabyLM baselines in the Strict-small and Strict tracks. Furthermore, we present the results from OLMo-1B. Finally, as a reference point, we include the results of children who were tested on the CVT by Surian et al. (1996). In the following, we break down the results by conversational maxim, model architecture and model size.

Results by Gricean Maxim. As shown in Table 2, model performance varies considerably across different maxims in both the Strict-small and Strict tracks. The maxims Quantity I (*Be informative*) and Quantity II (*Avoid redundant information*) are consistently the most challenging, with average accuracies for all BabyLMs peaking at only 0.59. In contrast, Quality (*Be truthful*) emerges as the easiest category for most BabyLMs with the average accuracy as high as 0.74. Relation (*Be relevant*) and Politeness (*Be polite*) generally fall in between these extremes (although exceptions apply), with the average accuracies above chance but below the best Quality results.

This pattern suggests that factuality is easier for BabyLMs to capture from limited data, likely because it can be learned from explicit statements and lexical associations in the training data, whereas judgments of informativeness and redundancy require more context-sensitive reasoning. The representative examples from the original CVT in Table 1 illustrate this: in Quantity I, while *"with milk"*

⁵Available at <https://github.com/kanishkamisra/minicons>.

falls in the range of pragmatically accepted answers, *"in a cup"* might contain tokens (or token combinations) that are more frequent in the training data. Across the other maxims, such frequent continuations may similarly override pragmatic appropriateness in model predictions.

To quantify how consistently models agree on these difficulty patterns, we ranked maxims per model and computed Kendall’s W . Agreement was high in the Strict-small track ($W = 0.80$) and moderate in the Strict track ($W = 0.68$), indicating that models trained on less data tend to exhibit more similar difficulty patterns. In both tracks, the maxims Quantity I and Quantity II were generally ranked as the hardest maxims. Full rankings and statistics are reported in Table 3.

Inter-maxim Correlations. To examine whether performance on different maxims co-varies across models, we computed Pearson correlations between per-model accuracies across maxims for each track (Figure 1). In the Strict-small track, correlations between maxims tend to be extreme. Quantity I and Quantity II show an almost perfect correlation ($r = 0.953$). The same track also reveals a striking near-perfect correlation between Relation and Politeness ($r = 0.999$), suggesting that topicality and politeness violations may be treated in similar ways. In contrast, Quality stands out in the Strict-small track as largely decoupled from the other maxims (near-zero or negative correlations), which may indicate that detecting literal implausibility (e.g., *"on the moon"*, Table 1) behaves independently of other pragmatic abilities under data constraints. In the Strict track, correlations are overall weaker and more varied, with some negative associations appearing for pairs that were strongly positive in the Strict-small track (e.g., Politeness vs. Quantity I, $r = -0.155$), possibly reflecting a partial decoupling of politeness from topicality and informativeness once models have more data. Notably, the maxim of Quality strongly correlates with Quantity I in the Strict track ($r = 0.931$), despite being among the easiest maxims in difficulty rankings (Table 3), showing that correlation patterns capture co-variation rather than absolute difficulty. Overall, the shift from rather extreme relations in the Strict-small track to more varied and generally weaker associations in the Strict track complements the prior results, suggesting that with more training data, models begin to differentiate more between pragmatic

	Quantity I	Quantity II	Quality	Relation	Politeness	Overall
Strict-small BabyLMs						
GPT-2	0.59	0.59	0.61	0.66	0.80	0.65
Baby Llama	0.56	0.49	0.74	0.63	0.68	0.62
LTG BERT	0.45	0.45	0.76	0.58	0.46	0.54
Roberta	0.33	0.36	0.64	0.60	0.57	0.50
<i>Average</i>	0.48	0.47	0.69*	0.61	0.63	
Strict BabyLMs						
GPT-2	0.60	0.68	0.76	0.72	0.76	0.70
Baby Llama	0.55	0.64	0.75	0.67	0.70	0.66
LTG BERT	0.64	0.57	0.79	0.74	0.52	0.65
Roberta	0.51	0.47	0.67	0.68	0.59	0.58
<i>Average</i>	0.58	0.59	0.74*	0.70	0.64	
LLM						
OLMo-1B	0.76	0.83	0.83	0.84	0.67	0.79
Children						
	0.58	0.78	1.0	1.0	0.93	0.86

Table 2: Accuracy scores across the Gricean maxims. *Strict-small* models are pretrained on <10M tokens and *Strict* models on <100M. OLMo-1B is pretrained on 3T tokens. All models are evaluated on 2,250 items, while child accuracy scores are from 8 neurotypical children from Surian et al. (1996). For the Strict-small and Strict groups, the highest score of each maxim is bolded. The highest average across all maxims is marked with (*).

Model	Quant. I	Quant. II	Qual.	Rel.	Polite
<i>Strict-small</i>					
GPT-2	●	●	●	●	●
Baby Llama	●	●	●	●	●
LTG BERT	●	●	●	●	●
RoBERTa	●	●	●	●	●
<i>Strict</i>					
GPT-2	●	●	●	●	●
Baby Llama	●	●	●	●	●
LTG BERT	●	●	●	●	●
RoBERTa	●	●	●	●	●

Table 3: Maxim difficulty rankings for each model (● = easiest, ● = hardest). Kendall’s W : Strict-small = 0.80, $\chi^2(4) = 12.80$, $p = 0.012$; Strict = 0.68, $\chi^2(4) = 10.80$, $p = 0.029$.

dimensions.

Effects of Architecture and Model size. Referring to Table 2, it is further notable that across both tracks, GPT-2 consistently outperforms LTG-BERT and RoBERTa models, with BabyLLaMA typically ranking second. RoBERTa BabyLM performs worst overall, particularly on the Quantity maxims. Increasing the training data from the Strict-small (<10M tokens) to the Strict (<100M tokens) track generally improves performance, especially for Quantity I and Quantity II, where average scores increase by 0.09–0.12. Gains for Quality and Relation are more modest, while Politeness scores remain similar across tracks. Interestingly, GPT-2’s Politeness score decreases in the Strict

track, suggesting that greater exposure to varied language might introduce alternative patterns that increase uncertainty in politeness judgments.

The large-scale OLMo 1B model substantially outperforms all BabyLM baselines across the maxims, with the exception of GPT-2 Strict-small, which scored 0.80 in Politeness. OLMo 1B scored substantially higher on the maxim of Relation compared to BabyLMs, indicating that sensitivity to topic relevance tends to emerge with increased training data.

The differences of scores may also reflect how auto-regressive and masked LMs handle conversational flow: in our setup, probability assignment to an entire answer benefits from modeling sequences as coherent continuations rather than token-masked completions. The performance gap between Strict-small and Strict models also indicates that increased training data helps, but does not eliminate the persistent difficulties with Quantity-related judgments.

Comparison to Child Performance. In Surian et al. (1996), neurotypical children showed an overall high accuracy (0.86), likely reflecting the development of Theory of Mind (Baron-Cohen et al., 1985), the impairment of which damages recognizing speaker’s intended meaning. BabyLMs share some similarities with children: like them, they perform best on the maxim of Quality and worst on the Quantity maxims (Table 2). However, they do not demonstrate the the same high performance




Figure 1: Pearson correlations between per-model accuracies across maxims. Top: Strict-small track. Bottom: Strict track.

as children in maxims of Relation and Politeness. Furthermore, children’s overall high accuracy indicates that, by school age, they are already highly sensitive to conversational norms. In contrast, BabyLMs reach only 0.50–0.70 of overall accuracy, underscoring a substantial gap in pragmatic competence between small-scale LMs and six-to-seven-year-old human speakers. The LLM shows a different pattern: it surpasses children on the Quantity maxims, indicating stronger performance on informativeness-related judgments, but still falls short on other maxims, especially Politeness, suggesting that socially grounded pragmatic norms do not emerge automatically from large-scale pre-training. Overall, both small-scale and large-scale models reveal persistent limitations in capturing the full range of conversational norms.

Summary and Discussion. Across both data tracks, the greatest deficits in BabyLM performance concerned judgments of informativeness

required for an appropriate response. The maxim of Quality was the easiest, with Relation and Politeness in between (Tables 2 and 3). Correlation analyses further revealed that, under severe data constraints, most maxims were strongly associated with at least one other (e.g., Quantity I and II; Relation and Politeness) (Figure 1). These associations tended to weaken with more data, indicating a shift toward more differentiated treatment of pragmatic categories. Among architectures, GPT-2 performed best overall, RoBERTa worst, and scaling from <10M to <100M tokens yielded the largest gains on Quantity, though sometimes at the expense of Politeness (Table 2), suggesting that autoregressive modeling and modest scaling benefit informativeness but may reduce social-pragmatic sensitivity. Compared to children, (0.86 overall, Table 2), BabyLMs mirrored the relative ordering of difficulty but scored substantially lower (0.50–0.70). The LLM (OLMo-1B) outperformed BabyLMs’ overall performance in all maxims and exceeded child performance on Quantity, yet lagged considerably on the remaining maxims, showing that large-scale pretraining enhances information-structuring abilities but offers limited gains in other dimensions of pragmatic understanding.

Our results align with earlier findings that pragmatic competence in language models scales with model size and training data but may remain below human levels. For instance, Hu et al. (2023) reported that GPT-2 with 117M parameters did not perform above chance when interpreting maxim-flouting utterances, in line with our observation that BabyLMs trained on <100M tokens perform rather poorly across maxims. At the other end of the scale, Yue et al. (2024) found that LLaMA 2 models with 13B parameters performed above chance across several Gricean maxims but still achieved only about half the human score, while GPT-4 matched human performance. In this context, our results with OLMo-1B suggest that large-scale pretraining can surpass child performance on informativeness but still leaves substantial gaps on more socially grounded maxims such as Politeness.

Overall, these findings indicate that scaling data and parameters improves some aspects of pragmatic reasoning in language models, while their absolute performance remains far from child-like. This underscores the importance of dedicated evaluation benchmarks targeting pragmatic abilities, ensuring that the developmental goals of the BabyLM challenge address this crucial aspect of language

use.

5 Conclusion

This paper introduced a novel large-scale benchmark for evaluating the pragmatic competence of language models, grounded in the Gricean maxims and adapted from a psycholinguistic test suite. Using 2,250 conversational items, we assessed BabyLM baseline models trained on constrained data alongside a large-scale 3T-token model and projected their performance on that of children. Our results indicate that increasing training data from <10M to <100M tokens leads to performance gains, yet BabyLMs remain below child-level competence. They demonstrated the lowest performance in assessing the appropriate amount of information for conversationally acceptable responses and did not exhibit a preference for answers that were more polite or contextually relevant. Furthermore, correlation analyses revealed that under data-limited conditions, models tend to conflate certain pragmatic competences, but these associations weaken with more data, suggesting that additional exposure allows models to more clearly differentiate between distinct pragmatic dimensions.

Our benchmark provides a linguistically grounded, scalable evaluation resource that enables systematic and comparable measurement of pragmatic behavior across models of different sizes and training regimes. By extending the BabyLM evaluation suite with a dedicated pragmatic benchmark, this work provides a tool for systematically tracking progress on this essential aspect of human-like language understanding.

Limitations and Future Directions. In this section, we state the limitations of our study and possible directions they offer for future work.

First, unlike other datasets (Zheng et al., 2021; Hu et al., 2023; Park et al., 2024), the conversational items in our dataset do not include detailed scenarios that are embedded before prompting models with dialogues. In certain contexts, the maxim-violator responses in our dataset could be in fact appropriate; for instance, the answer *"My best friend is Peter. He wears clothes."* (Table 1) would not be redundant in a scenario where others are unclothed. However, even within a minimal-context setting like the one implemented in this study, non-linguist participants have been shown to consistently favor responses that adhere to conversational maxims. For example, Okanda et al. (2015) applied

a revised Japanese version of the CVT and found that adults were able to identify non-cooperative answers and articulate the reasoning behind their judgments. Nevertheless, future work could expand our conversational items by introducing explicit context that would render maxim-violator responses cooperative (as in the above-mentioned example) to examine whether such framing would change model preferences.

Second, we acknowledge that model probability assignments may be influenced by the distributional properties of tokens independent of context, which can make evaluations based on sequence scores prone to bias. Future work could address this by expanding the dataset to include a wider range of lexical variations.

Third, we selected the 1B-parameter OLMo model due to computational constraints. Although our primary focus was on BabyLMs rather than large models, evaluating systems of varying sizes offers valuable comparative insights. Furthermore, as our results suggest that model architecture affects pragmatic performance, future work could test whether these patterns hold for other architectures, such as instruction-tuned or multimodal (Baby) models.

Finally, our dataset is currently limited to English; therefore, extending this evaluation to multilingual settings would allow for more robust conclusions and enable meaningful cross-linguistic comparisons of pragmatic competence.

Ethical Statement

Our work uses publicly available data from a psycholinguistic study on children; we do not conduct any new experiments involving human subjects. No personal information from the original study is shared, except for the mean age of participants as reported by the authors. We do not train any new models; instead, our evaluation dataset was automatically generated and subsequently reviewed and refined by two of the authors. The dataset consists of short, child-level conversational exchanges, and we believe that none of the items raise ethical concerns or reinforce biases toward sensitive groups. Our evaluation focuses exclusively on the pragmatic competence of language models. We do not address other potential harms or limitations of these systems such as discrimination, toxicity and misinformation Weidinger et al. (2021) which remain important areas for continued investigation and responsible deployment. We share our evalu-

ation dataset, code and model results publicly to facilitate future use and promote transparency.

Acknowledgements

We warmly thank the anonymous reviewers for their insightful comments. We acknowledge support from several projects and funding institutions: 1) “SAIL: SustAInable Life-cycle of Intelligent Socio-Technical Systems” (Grant ID NW21-059A), an initiative of the Ministry of Culture and Science of the State of Northrhine Westphalia; 2) Erasmus+, a European Union funding that facilitates periods of study or training abroad, 3) Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – CRC-1646, project number 512393437, project B02.

References

- Brian P Ackerman. 1981. When is a question not answered? the understanding of young children of utterances violating or conforming to the rules of conversational sequencing. *Journal of Experimental Child Psychology*, 31(3):487–507.
- Simon Baron-Cohen, Alan M. Leslie, and Uta Frith. 1985. Does the autistic child have a “theory of mind”? *Cognition*, 21(1):37–46.
- Ye-eun Cho and Seong mook Kim. 2024. [Pragmatic inference of scalar implicature by LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 10–20, Bangkok, Thailand. Association for Computational Linguistics.
- Leshem Choshen, Ryan Cotterell, Michael Y. Hu, Tal Linzen, Aaron Mueller, Candace Ross, Alex Warstadt, Ethan Wilcox, Adina Williams, and Chengxu Zhuang. 2024. [\[call for papers\] the 2nd babylm challenge: Sample-efficient pretraining on a developmentally plausible corpus](#). *Preprint*, arXiv:2404.06214.
- Daniel J Conti and Linda A Camras. 1984. Children’s understanding of conversational principles. *Journal of Experimental Child Psychology*, 38(3):456–463.
- Judith Degen and Michael Tanenhaus. 2014. [Processing scalar implicature: A constraint-based approach](#). *Cognitive Science*.
- Ryan Doran, Gregory Ward, Meredith Larson, Yaron McNabb, and Rachel E Baker. 2012. A novel experimental paradigm for distinguishing between what is said and what is implicated. *Language*, 88(1):124–154.
- Allyson Ettinger. 2020. [What bert is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. 2021. [Causal analysis of syntactic agreement mechanisms in neural language models](#). 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 1828–1843, Online. Association for Computational Linguistics.
- Daniel Fried, Nicholas Tomlin, Jennifer Hu, Roma Patel, and Aida Nematzadeh. 2023. [Pragmatics in language grounding: Phenomena, tasks, and modeling approaches](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12619–12640, Singapore. Association for Computational Linguistics.
- Kanishk Gandhi, J.-Philipp Fränken, Tobias Gerstenberg, and Noah D. Goodman. 2024. Understanding social reasoning in language models with language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Herbert P Grice. 1975. Logic and conversation. In *Speech acts*, pages 41–58. Brill.
- Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, and 24 others. 2024. [OLMo: Accelerating the science of language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand. Association for Computational Linguistics.
- Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko, and Edward Gibson. 2023. [A fine-grained comparison of pragmatic language understanding in humans and language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4194–4213, Toronto, Canada. Association for Computational Linguistics.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. [A systematic assessment of syntactic generalization in neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Sophie Jentsch and Kristian Kersting. 2023. [ChatGPT is fun, but it is not funny! humor is still challenging large language models](#). In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 325–340, Toronto, Canada. Association for Computational Linguistics.

- Simeon Junker, Manar Ali, Larissa Koch, Sina Zarrieß, and Hendrik Buschmeier. 2025. [Are multimodal large language models pragmatically competent listeners in simple reference resolution tasks?](#) In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24101–24109, Vienna, Austria. Association for Computational Linguistics.
- Jad Kabbara and Jackie Chi Kit Cheung. 2022. Investigating the performance of Transformer-Based NLI models on presuppositional inferences. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 779–785, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Florian Kankowski, Torgrim Solstad, Sina Zarrieß, and Oliver Bott. 2025. [Instruction tuning modulates discourse biases in language models.](#)
- Carina Kauf and Anna Ivanova. 2023. [A better way to do masked language model scoring.](#) *Preprint*, arXiv:2305.10588.
- Daria Kryvosheieva and Roger Levy. 2025. [Controlled evaluation of syntactic knowledge in multilingual language models.](#) In *Proceedings of the First Workshop on Language Models for Low-Resource Languages*, pages 402–413, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Charlotte Kurch, Margarita Ryzhova, and Vera Demberg. 2024. [Large language models fail to derive atypicality inferences in a human-like manner.](#) In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 86–100, Bangkok, Thailand. Association for Computational Linguistics.
- Clara Lachenmaier, Judith Sieker, and Sina Zarrieß. 2025. [Can LLMs ground when they \(don’t\) know: A study on direct and loaded political questions.](#) In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14956–14975, Vienna, Austria. Association for Computational Linguistics.
- Andrew Lampinen. 2024. [Can language models handle recursively nested grammatical structures? a case study on comparing models and humans.](#) *Computational Linguistics*, 50(3):1441–1476.
- 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.](#) *Preprint*, arXiv:1907.11692.
- Bolei Ma, Yuting Li, Wei Zhou, Ziwei Gong, Yang Janet Liu, Katja Jasinskaja, Annemarie Friedrich, Julia Hirschberg, Frauke Kreuter, and Barbara Plank. 2025a. [Pragmatics in the era of large language models: A survey on datasets, evaluation, opportunities and challenges.](#) In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8679–8696, Vienna, Austria. Association for Computational Linguistics.
- Ziqiao Ma, Jing Ding, Xuejun Zhang, Dezhi Luo, Ji-ahé Ding, Sihan Xu, Yuchen Huang, Run Peng, and Joyce Chai. 2025b. [Vision-language models are not pragmatically competent in referring expression generation.](#) *Preprint*, arXiv:2504.16060.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Kanishka Misra. 2022. [minicons: Enabling flexible behavioral and representational analyses of transformer language models.](#) *Preprint*, arXiv:2203.13112.
- Mako Okanda, Kosuke Asada, Yusuke Moriguchi, and Shoji Itakura. 2015. Understanding violations of gricean maxims in preschoolers and adults. *Frontiers in psychology*, 6:901.
- OpenAI. 2024. [Gpt-4 technical report.](#) *Preprint*, arXiv:2303.08774.
- Francesca Panzeri and Francesca Foppolo. 2021. Children’s and adults’ sensitivity to gricean maxims and to the maximize presupposition principle. *Frontiers in Psychology*, 12:624628.
- Dojun Park, Jiwoo Lee, Hyeyun Jeong, Seohyun Park, and Sungeun Lee. 2024. [Pragmatic competence evaluation of large language models for the Korean language.](#) In *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation*, pages 256–266, Tokyo, Japan. Tokyo University of Foreign Studies.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Giulia Rambelli, Emmanuele Chersoni, Claudia Colacciani, and Marianna Bolognesi. 2024. [Can large language models interpret noun-noun compounds? a linguistically-motivated study on lexicalized and novel compounds.](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11823–11835, Bangkok, Thailand. Association for Computational Linguistics.
- Jacopo Romoli and Florian Schwarz. 2015. [An Experimental Comparison Between Presuppositions and Indirect Scalar Implicatures](#), pages 215–240. Springer International Publishing, Cham.
- David Samuel, Andrey Kutuzov, Lilja Øvrelid, and Erik Velldal. 2023. [Trained on 100 million words and still in shape: BERT meets British National Corpus.](#) In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1954–1974, Dubrovnik, Croatia. Association for Computational Linguistics.

- Elisa Sanchez-Bayona and Rodrigo Agerri. 2025. [Metaphor and large language models: When surface features matter more than deep understanding](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 17462–17477, Vienna, Austria. Association for Computational Linguistics.
- Natalie Shapira, Guy Zwirn, and Yoav Goldberg. 2023. How well do large language models perform on faux pas tests? In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10438–10451, Toronto, Canada. Association for Computational Linguistics.
- Judith Sieker, Oliver Bott, Torgrim Solstad, and Sina Zarriß. 2023. [Beyond the bias: Unveiling the quality of implicit causality prompt continuations in language models](#). In *Proceedings of the 16th International Natural Language Generation Conference*, pages 206–220, Prague, Czechia. Association for Computational Linguistics.
- Judith Sieker, Clara Lachenmaier, and Sina Zarriß. 2025. [LLMs struggle to reject false presuppositions when misinformation stakes are high](#). *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47.
- Judith Sieker and Sina Zarriß. 2023. [When your language model cannot Even do determiners right: Probing for anti-presuppositions and the maximize presupposition! principle](#). In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 180–198, Singapore. Association for Computational Linguistics.
- Luca Surian, Simon Baron-Cohen, and Heather van der Lely. 1996. [Are children with autism deaf to gricean maxims?](#) *Cognitive neuropsychiatry*, 1:55–72.
- Luca Surian, Mariantonia Tedoldi, and Michael Siegal. 2010. Sensitivity to conversational maxims in deaf and hearing children. *Journal of Child Language*, 37(4):929–943.
- Lyn Tieu, Jacopo Romoli, Peng Zhou, and Stephen Crain. 2015. [Children’s knowledge of free choice inferences and scalar implicatures](#). *Journal of Semantics*, 33(2):269–298.
- Inar Timiryasov and Jean-Loup Tastet. 2023. [Baby llama: knowledge distillation from an ensemble of teachers trained on a small dataset with no performance penalty](#). *Preprint*, arXiv:2308.02019.
- Sean Trott, Cameron Jones, Tyler Chang, James Michaelov, and Benjamin Bergen. 2023. [Do large language models know what humans know?](#) *Cognitive Science*, 47(7):e13309.
- Polina Tsvilodub, Paul Marty, Sonia Ramotowska, Jacopo Romoli, and Michael Franke. 2024. [Experimental pragmatics with machines: Testing llm predictions for the inferences of plain and embedded disjunctions](#). *Preprint*, arXiv:2405.05776.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). *Preprint*, arXiv:1804.07461.
- Alex Warstadt, Leshem Choshen, Aaron Mueller, Adina Williams, Ethan Wilcox, and Chengxu Zhuang. 2023. [Call for papers – the babylm challenge: Sample-efficient pretraining on a developmentally plausible corpus](#). *Preprint*, arXiv:2301.11796.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, and 1 others. 2021. [Ethical and social risks of harm from language models](#). *arXiv preprint arXiv:2112.04359*.
- Shisen Yue, Siyuan Song, Xinyuan Cheng, and Hai Hu. 2024. [Do large language models understand conversational implicature- a case study with a Chinese sitcom](#). In *Proceedings of the 23rd Chinese National Conference on Computational Linguistics (Volume 1: Main Conference)*, pages 1270–1285, Taiyuan, China. Chinese Information Processing Society of China.
- Wenting Zhao, Justin Chiu, Claire Cardie, and Alexander Rush. 2023. [Abductive commonsense reasoning exploiting mutually exclusive explanations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14883–14896, Toronto, Canada. Association for Computational Linguistics.
- Zilong Zheng, Shuwen Qiu, Lifeng Fan, Yixin Zhu, and Song-Chun Zhu. 2021. [GRICE: A grammar-based dataset for recovering implicature and conversational rEasoning](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2074–2085, Online. Association for Computational Linguistics.

A Appendix

A.1 Dataset Augmentation Details

In order to expand CVT for model evaluation, we employed the GPT-4 chat interface (OpenAI, 2024) to synthesize similar items. First, we provided GPT-4 with a brief description of the maxims and examples from the CVT, and asked it to generate 25 corresponding dialogue templates, each accompanied by two paraphrased versions of the question. Answers were unchanged in the paraphrased versions; this was due to some answers being too short to paraphrase, and we chose a unified method for all items. Next, each template was presented with its original CVT example, and GPT-4 was asked to produce four new conversations per template, using vocabulary appropriate for children. Finally, the paraphrased versions of both the original CVT items and the newly created conversations were generated based on their respective templates. This resulted in a dataset consisting of 25 CVT items and 350 GPT-generated ones. Table 4 shows a concrete example of this pipeline.

Two of the authors manually reviewed all 375 items and made adjustments based on the following criteria:

- The follower’s answer does not follow the maxim.
- The violator’s answer does not violate the maxim.
- The paraphrased versions do not correctly paraphrase the questions.
- The words exceed child-level vocabulary.
- The questions and/or answers are unnatural or ungrammatical.
- The answers are inadequate for model evaluation.

The last criterion reflects our effort to minimize superficial differences between the follower and violator answers wherever possible. Since our evaluation compares model probabilities for each answer pair, irrelevant lexical differences would distort the results. For example, in response to the question *Who is your best friend?*, if one answer was *My best friend is **John**. He goes to my school.* and the other was *My best friend is **Peter**. He wears clothes.*, not replacing the names with a single name would introduce noise unrelated to pragmatic reasoning. However, in some cases, such standardizations were not applicable due to the nature of the maxims being

tested; for example, in the case of the maxim of Relation, altering the content is necessary for the violation.

Once approved, we diversified the dataset by assigning speaker names (Leslie, Joan, and Thomas) and rotating them such that each conversation had one name as the questioner and another as the responder, used consistently across both responses. The names were randomly selected as English common names from the firstname database at <https://github.com/KarlAmort/firstname-database>. This process yielded a final dataset of 2,250 conversational items, since six variations with names were generated from each of the 375 curated conversations.

Stage	Example
1. GPT-4 is prompted with a dialogue from CVT	Question: <i>What would you like to buy in this shoe-shop?</i> Follower: <i>A pair of trainers.</i> Violator: <i>A pair of shoes.</i>
2. GPT-4 generates a template from the dialogue	Question: <i>What would you like to buy in this [type of store]?</i> Follower (specific, informative): [A concrete item you might buy in that specific store] Violator (vague, unhelpful): [A general category that lacks specific information]
3. GPT-4 generates four new dialogues from the template of stage 2	1. Question: <i>What would you like to buy in this toy shop?</i> Follower: <i>A teddy bear.</i> Violator: <i>A toy.</i> 2. Question: <i>What would you like to buy in this bakery?</i> Follower: <i>A chocolate muffin.</i> Violator: <i>Something to eat.</i> 3. Question: <i>What would you like to buy in this bookshop?</i> Follower: <i>A Harry Potter book.</i> Violator: <i>A book.</i> 4. Question: <i>What would you like to buy in this clothes shop?</i> Follower: <i>A red jacket.</i> Violator: <i>Some clothes.</i>
4. GPT-4 generates 2 paraphrased questions of the template of stage 2	1. <i>Is there something you'd like to get from this [type of store]?</i> 2. <i>What are you looking for in this [type of store]?</i>
5. GPT-4 generates 2 paraphrased versions of the original CVT dialogue given in stage 1 and all other dialogues generated in stage 3	1. Question: <i>Is there something you'd like to get from this shoe-shop?</i> Follower: <i>A pair of trainers.</i> Violator: <i>A pair of shoes.</i> 2. Question: <i>What are you looking for in this shoe-shop?</i> Follower: <i>A pair of trainers.</i> Violator: <i>A pair of shoes.</i> And so on

Table 4: An example of the data augmentation pipeline for the maxim of Quantity I (*Be informative*). The Follower adheres to the maxim, while the Violator does not. For each dialogue in CVT, four more examples with child-level vocabulary were created. Later, two paraphrased versions for all dialogues (those from CVT and GPT-generated ones) were synthesized and added.

A.2 Dataset Examples

We depict a few examples from the dataset. The maxim-violator’s answer is marked with (*).

Maxim of Quantity I (Be informative):

- Leslie: What did you eat for supper?
Thomas: Tomato soup.
Thomas: A dish.*
- Leslie: What did you see at the zoo?
Joan: The lions.
Joan: Some animals.*
- Leslie: What did you get for Christmas?
Joan: A gift.*
Joan: A toy train.
- Joan: How do you prefer your pancakes?
Leslie: On a plate.*
Leslie: With maple syrup.

Maxim of Quantity II (Avoid redundant information):

- Leslie: Who is your neighbor?
Thomas: My neighbor is Mr. Tom. He has a dog.
Thomas: My neighbor is Mr. Tom. He lives in a house.*
- Joan: What pet do you like?
Leslie: I like puppies and kittens which are pets.*
Leslie: I like puppies and kittens which are cute.
- Leslie: Where did you go last weekend?
Joan: I went to grandma’s house and baked cookies.
Joan: I went to grandma’s house and I didn’t stay in my room.*
- Joan: Which is your favourite fruit?
Thomas: Watermelon which is a fruit.*

Thomas: Watermelon which is juicy.

Maxim of Quality (Be truthful):

- Thomas: Is there any more popcorn?
Leslie: Yes, there's a bowl in the kitchen.
Leslie: Yes, it's raining popcorn outside.*
- Leslie: Where do you do your homework?
Joan: I do them on a dragon's back.*
Joan: I do them in my room.
- Joan: Do you have any pets?
Leslie: Yes, I have a cat and a fish.
Leslie: Yes, I have a thousand elephants.*
- Leslie: Why don't you come outside?
Thomas: Because I'm helping mom bake.
Thomas: Because I'm on a spaceship.*

Maxim of Relation (Be relevant):

- Joan: What did you do on the weekend?
Thomas: I went to the zoo.
Thomas: My socks are green.*
- Thomas: What do you like to play?
Leslie: I like to play tag.
Leslie: I like chocolate cake.*
- Joan: What is your favourite animal?
Thomas: I like pencils best.*
Thomas: I like pandas best.
- Thomas: What songs do you know?
Joan: I know "Twinkle Twinkle Little Star."
Joan: I know how to tie my shoes.*

Maxim of Politeness (Be polite):

- Thomas: Do you like my new haircut?
Joan: It looks awful.*
Joan: It looks nice.
- Thomas: Would you like to try some of my cake?
Joan: No, thanks.
Joan: No, it's disgusting.*
- Joan: May I use your calculator?
Thomas: No, don't touch my stuff.*
Thomas: No, sorry, I need it right now, but you can use it after.
- Leslie: Could you help me with my puzzle?
Thomas: Do it by yourself.*
Thomas: Sure, after I finish this one.