Large Language Models and the Argument From the Poverty of the Stimulus

Nur Lan, Emmanuel Chemla, and Roni Katzir

November 2023

Abstract

How much of our linguistic knowledge is innate? According to much of theoretical linguistics, a fair amount. One of the best-known (and most contested) kinds of evidence for a large innate endowment is the so-called *argument from the poverty of the stimulus* (APS). In a nutshell, an APS obtains when human learners systematically make inductive leaps that are not warranted by the linguistic evidence. A weakness of the APS has been that it is very hard to assess what is warranted by the linguistic evidence. Current Artificial Neural Networks appear to offer a handle on this challenge. Wilcox et al. (2022) use such models to examine the available evidence as it pertains to wh-movement. They conclude that the (presumably linguistically-neutral) networks acquire an adequate knowledge of wh-movement, thus undermining an APS in this domain. We examine the evidence further and show that the networks do not, in fact, succeed in acquiring wh-movement. More tentatively, our findings suggest that the failure of the networks is due to the insufficient richness of the linguistic input, thus supporting the APS.

1 Background: innateness and the argument from the poverty of the stimulus

One way in which linguists have argued that humans are born with nontrivial biases is through cases where speakers' linguistic knowledge goes beyond what seems warranted by the data they were exposed to. If humans systematically arrive at this knowledge given the data while linguistically-neutral learners do not, then humans are not linguistically neutral: they come to the task of language acquisition prepared. Reasoning of this kind is known as an *argument from the poverty of the stimulus* (APS), and since its introduction by Noam Chomsky over 50 years ago it has been central to the study of the human linguistic

capacity.[1,2] Here we will focus on one APS, concerning wh-movement, but various other APSs have been discussed in the literature based on a range of empirical phenomena such as *one*-substitution (introduced in Baker 1978), subject-auxiliary inversion (introduced in Chomsky 1971), and plurals within compounds (introduced in Gordon 1985).

The APSs just mentioned (and others like them) have been taken to argue for nontrivial innate biases in humans. For example, the APS from subject-auxiliary inversion has been taken to support an innate bias for hierarchical transformations over linear ones. The APS from wh-movement that we discuss below will similarly support an intricate bias that a linguistically-neutral learner is not expected to have. The same holds for other APSs in the literature. In this, these APSs go beyond the early observation that children can produce and understand unboundedly many sentences after encountering only a finite number of sentences (Chomsky 1957, p. 15). While generalizing from a finite input to an infinite language is perhaps not entirely trivial, it is something that most learning algorithms do. And importantly, this ability does not imply any biases that a linguistically-neutral will not have.

While the APS has been central to linguistic reasoning, it has also generated much controversy. Contesting a given APS requires challenging either the knowledge attained by humans or the information available to the child learner. It is the latter that often comes under attack. The reason for this vulnerability is that it is extremely difficult to assess what information exactly is available to the child over the relevant time period (often years of exposure) and hard to tell what a general-purpose, linguistically-neutral learner would do with this kind of information. One can try to look for pieces of evidence that seem relevant for the knowledge at stake — e.g., as done for the case of subject-auxiliary inversion in English by Legate and Yang (2002) — but as noted by Perfors et al. (2011), this methodology runs the risk of underestimating the available information: even if we fail to find the evidence we are looking for, a general-purpose learner might be able to take advantage of other sources of information. This methodology also risks *over*estimating the available information: even if we find several instances of the evidence we are after, a general-purpose learner might treat those instances as noise and fail to draw the inference that we intuitively expect it to. In the absence of an actual learner that can use the information that is available in an entire corpus it is just very hard to estimate whether the data support the knowledge

---

[1]The general considerations behind the APS are discussed already in chapter 1 of Chomsky 1965. Further considerations are discussed in Chomsky 1971, pp. 26–8, Chomsky 1975, pp. 30ff., and Chomsky, 1980, pp. 42ff., as well as in much subsequent work.

In addition to the APS, linguists have also identified other sources of evidence supporting the innateness of nontrivial linguistic knowledge. For example, there are arguments from the *richness* of the stimulus, where a pattern that is clearly represented in the input data and would be easily picked up by a linguistically-neutral learner is simply ignored by human learners. Evidence from typological asymmetries has also played a very important role in linguistic reasoning. A proper discussion of such sources of evidence falls outside the scope of the present paper, and in what follows we focus exclusively on the APS.

[2]Throughout the discussion we set aside the question of whether the knowledge under consideration is specific to linguistics (and, if so, how much of it is purely syntactic) or whether it is shared with other cognitive domains. Our sole focus is on whether a neutral learner would be justified in acquiring the relevant knowledge based on a given linguistic input.

under consideration.[3]

How then can we reason about the information available to the child and ask whether it suffices to support the acquisition of a given piece of knowledge by a linguistically-neutral learner? In an ideal world, one would (a) take a sufficiently powerful learner that can be seen to not be biased in favor of the relevant knowledge; (b) train this learner on a corpus that corresponds to the linguistic input that children receive; and (c) check whether the learner has indeed acquired the knowledge under consideration. In such an ideal world, one might perhaps be able to work with a Bayesian program induction algorithm for a general-purpose programming language such as LISP or Python using a description-length prior, where neither the programming language nor the learning algorithm can be taken to bias the learning in the direction of human-like knowledge of wh-movement, and where the knowledge acquired by the algorithm can be directly inspected at stage (c).

In the actual world, combining (a) through (c) is currently impossible. For many years, the combination of (a) and (b) was already a major barrier, since many learning models could not be trained on realistic corpora (corresponding to years of human linguistic experience) in the first place, while those models that could be trained on such corpora — such as n-gram models — were insufficiently powerful to acquire or even represent sophisticated linguistic knowledge such as wh-movement dependencies.

The challenge of assessing the information available to the child has become less of an obstacle lately, with the advent of Large Language Models (LLMs). These models, which rely on modern architectures of artificial neural networks (ANNs), can be trained on very large corpora and are generally quite successful in acquiring sequential dependencies. Moreover, they can be tentatively taken to be linguistically neutral. The qualification of this last point is due to how poorly the actual biases of current LLMs are understood: until more is known about these biases, and as correctly cautioned by Rawski and Heinz (2019), any claims about the neutrality of these models must be taken as tentative. Still, it strikes us as reasonable to assume that current LLMs are not particularly biased *against* the linguistic dependencies under consideration. This is especially so since these models have evolved over the past decades so as to succeed in capturing key patterns in linguistic sequences; therefore, if they do have linguistically-relevant biases after all, those are likelier to be in favor of the patterns under consideration rather than against them. Consequently, if the models fail to acquire an adequate approximation of the relevant dependencies, this failure

---

[3]See Pullum and Scholz (2002), Lidz et al. (2003), Foraker et al. (2009), Hsu and Chater (2010), Berwick et al. (2011), Perfors et al. (2011), and Pearl and Sprouse (2013), among others, for much relevant discussion.

In studies of analogous inductive leaps in other species, this worry regarding the input has been addressed by controlling the information available to the learners (see, e.g., Dyer and Dickinson 1994). To a certain extent this can be done with humans in experiments of artificial-grammar learning (see, e.g., Wilson 2006). But for the main APSs in the literature, which concern the normal course of child language acquisition, controlling the information available to the learner is not an option.

can be taken to be informative.[4] To some extent, then, LLMs address both (a) and (b) above.[5]

Unfortunately, LLMs still pose a challenge with respect to (c): it is all but impossible to inspect them and determine what they know. Nevertheless, in recent work, Wilcox et al. (2022; WFL) make the case for still using LLMs to evaluate the APS. Building on a methodological paradigm from psycholinguistics (and as discussed in more detail below), WFL use LLMs to study wh-movement and argue that current models acquire an adequate knowledge of wh-movement from their respective training corpora. This, WFL suggest, debunks an APS: one that says that the input is insufficiently rich to allow a general-purpose learner to acquire islands.[6] More broadly, WFL take their results to suggest that the innate biases that humans bring to the task of language acquisition might be linguistically neutral.

The present paper extends WFL's probing of LLMs' knowledge of wh-movement, arriving at conclusions that are at odds with those of WFL. We start, in Section 2, with a brief overview of the general setup for the rest of the paper. Among other things we discuss how LLMs can be used as tools for assessing the information in a given corpus without assuming that these models are cognitively plausible in any way. In Section 3 we discuss the success of LLMs in simple cases of wh-dependencies, as noted by WFL. In section 4 we show that the scope of the LLMs' success is rather limited. In particular, LLMs fail to adequately approximate human knowledge of a much-studied family of cases, falling under the labels of parasitic gaps and across-the-board movement, in which certain additional gaps make an otherwise problematic gap inside an island acceptable. It is cases such as these that are typically taken by linguists to suggest an APS, and our findings show that the performance of the LLMs does not, in fact, debunk this APS. Moreover, the failure of the LLMs suggests, albeit weakly, that an input of the kind that children see is not rich enough to support the acquisition of wh-movement by a linguistically-neutral learner. In other words, the failure of the LLMs offers preliminary support for the APS. In section 5, we strengthen this support by considering several alternative explanations for the failure of the LLMs that attribute the failure to weaknesses of the ANNs on which these LLMs are based rather than to the poverty of the stimulus. In particular, we consider the possibility that the ANNs fail because of specific lexical choices in the test data (section 5.1), which we address by varying the lexical choices; that they fail because of a general preference

---

[4]As noted in the literature, there are cases where access to world knowledge can give human children an unfair advantage over ANNs, in which case the failure of ANNs to acquire a given pattern might not be indicative. See Warstadt and Bowman (2022) for discussion. In the present case we do not see how world knowledge might assist in acquiring the properties of wh-movement that we focus on.

[5]Long before the current models, earlier ANN architectures were used in debates of the APS, and in particular in attempts to argue against various versions of it (see Elman et al. 1996, Lewis and Elman 2001, and Reali and Christiansen 2005, and see Berwick et al. 2011 for a critical analysis of some earlier attempts). Early ANNs, however, were limited in their capacities and generally trained on small corpora, and it is unclear whether they could be used to reason about whether a corpus that roughly corresponds to children's linguistic exposure supports the acquisition of complex grammatical knowledge. The ability of current models to train on realistically large corpora is a helpful step towards using them constructively in debates about the APS.

[6]See Pearl and Sprouse (2013) and Phillips (2013) for earlier discussion of APS in the context of acquiring islands.

for non-gap continuations (section 5.2), which we address, following WFL, by comparing the preference for a gap when a wh-element is present upstream and when it is absent; and that they fail because they are inherently incapable of approximating the patterns under consideration, which we address by retraining one of the models on corpora that are clearly *not* impoverished with respect to the relevant patterns (section 5.3). In all cases we provide evidence against the relevant explanations, thus strengthening the APS. Section 6 concludes.

2 The general setup

Simplifying considerably, a *gap*, such as the missing complement of 'with' in (1a) and (1c), appears if and only if it is preceded by an appropriate *filler*, such as the wh-phrase 'who' in (1a) and (1b). When there is both a filler and a gap (1a) or neither (1d) the result is good; when there is a filler and no gap (1b) or a gap and no filler (1c) the result is bad.[7]

(1)   a.   I know <u>who</u> you talked with __ yesterday. $(+filler,+gap)$

b.   * I know <u>who</u> you talked with Mary yesterday. $(+filler,-gap)$

c.   * I know <u>that</u> you talked with __ yesterday. $(-filler,+gap)$

d.   I know <u>that</u> you talked with Mary yesterday. $(-filler,-gap)$

There is much further nuance to wh-movement, some of which we will briefly mention below. For now, let us consider how one might check if the input data are rich enough for a linguistically-neutral learner to acquire the knowledge of wh-movement.

As mentioned, LLMs are very opaque. In particular, we cannot at present check whether they believe that a given continuation such as 'yesterday' or 'Mary' is grammatical following a given prefix such as 'I know who/that you talked about'. In fact, it is not clear whether current models even have a notion of grammaticality to begin with.

What LLMs do tell us is how *likely* they consider any given continuation. The problem is that grammaticality and probability are generally very different notions. And while the two are correlated — many ungrammatical continuations are also unlikely on any sensible notion of probability, and grammatical continuations are sometimes probable — this correlation is far from perfect (see Chomsky 1957, Berwick 2018, and Sprouse et al. 2018, among others, for relevant discussion). In particular, many grammatical continuations are highly unlikely; e.g., 'splat' is a grammatical but unlikely continuation of 'John would like to eat a freshly-made'. And in some cases an ungrammatical continuation can be likely; e.g., 'is' is a likely but ungrammatical continuation of 'The keys to the cabinet', an instance

---

[7]In order to make it possible to alternate the $\pm filler$ condition, and following WFL, we embed the relevant examples under 'I know': 'I know <u>who</u>...' $(+filler)$ vs. 'I know <u>that</u>...' $(-filler)$.

of so-called *agreement attraction*.[8]

In the cases we are interested in here, however, probability and grammaticality are often quite well aligned, and it is easy to find examples such as (1) in which the grammatical continuation is significantly more probable than the ungrammatical one on any sensible notion of probability. So if we focus on such cases where grammaticality and probability are aligned, and if ANNs are sufficiently good learning models — at least, good enough to provide a crude approximation of the pattern under consideration — then we can use the probabilistic predictions of the resulting LLMs to evaluate the APS. If the LLMs systematically assign a much higher probability to the grammatical continuation, this can be taken to suggest that the pattern of wh-movement is represented in the input data sufficiently well so as to make it worth a learner's while to at least acquire an approximation of the pattern for simple cases. While it remains unclear, as mentioned above, whether ANNs themselves have a representation of grammaticality as distinct from probability or whether they can learn or even represent the true pattern, their success would suggest that a good linguistically-neutral learner that does have such representational abilities might acquire the pattern. Conversely, if the LLMs do not systematically assign a much higher probability to the grammatical continuation, this would indicate that the pattern of wh-movement is not sufficiently well represented in the input data to merit the acquisition of at least a passable approximation of it by a linguistically-neutral learner; in other words, the APS holds. In this way, LLMs — even if their representational inadequacies prevent them from providing more than a crude approximation of the pattern under consideration — can serve as useful proxies for future general-purpose learners and help us reason about the information available in the input data.

In section 5 below we discuss residual worries with LLMs, even when used solely as proxies for better models and without making commitments to their representations and knowledge. First, however, let us examine what happens when several different LLMs are used in the domain of wh-movement.

3   LLMs succeed in very simple cases of wh-movement

How rich is the input, then, when it comes to filler-gap dependencies of the wh kind? WFL evaluate this question using four language models that achieved state-of-the-art results on various NLP benchmarks. Following them we also approach this question using a battery of language models. The models we use in our evaluation are the following, also summarized in Table 1: two models from Yedetore et al. (2023), an LSTM and a Transformer, that

---

[8]Agreement attraction is a performance error. Speakers make such errors when distracted or in a hurry but less so when given more time. ANNs do not make this distinction: when they give a higher probability to an ungrammatical continuation their response reflects a faulty knowledge rather than a resource problem. This serves to further illustrate the inadequacy of ANNs as models of linguistic cognition. However, since our goal is not to study ANNs for their own sake but rather to use them to evaluate the informativeness of the input data, the inadequacy of ANNs as cognitive models is not inherently problematic.

were trained on the CHILDES corpus of child-directed speech (MacWhinney, 2014);[9,10] an LSTM trained on English Wikipedia (Gulordava et al., 2018); a Transformer trained on English Wikipedia;[11] Open AI's GPT-2 (Radford et al., 2019); GPT-j (Wang and Komatsuzaki, 2021); and OpenAI's GPT-3 (Brown et al., 2020).[12] The two CHILDES-trained models, as mentioned, are taken from Yedetore et al. (2023). The LSTM trained on English Wikipedia and both GPT-2 and GPT-3 are used by WFL in their evaluation.[13]

In order to get a very rough sense of the number of years of linguistic experience that a given training corpus corresponds to we follow common practice (used also by WFL) based on Hart and Risley (1992)'s estimates about the number of words that American children typically hear during acquisition. According to these estimates, the models just mentioned were exposed to amounts of data ranging from ten months of linguistic experience (CHILDES LSTM and Transformer) through eight years of linguistic experience (Wikipedia LSTM and Transformer) to between 10 and 500 human lifetimes (GPT-2, GPT-3, and GPT-j); see Table 1. Indeed, WFL admit that the linguistic experience of some of the models is probably above and beyond that of children, and could thus weaken their argument against APS in case of successful learning by the models. However, in the case of a negative result, as in the current work, a large training corpus strengthens the argument: if these models are exposed to amounts of data that go beyond what children are exposed to and still don't learn the constructions under consideration, this serves to strengthen the APS for these phenomena.

In very simple cases such as (1) above, the LLMs considered by WFL assign a higher probability to the grammatical continuation than to the ungrammatical one. Above we mentioned that success in cases such as those considered here, where probability and ac-

---

[9]The models in Yedetore et al. (2023) were trained on utterances of 52 children between the ages of six months to 12 years. The total training size amounts to 9.6 million words. While quantitatively this does not amount to the input available throughout actual childhood, here we inherit the assumption in this work that training models on such data approximates to some extent the kind of input available to a child learner.

[10]Out of ten models per architecture (LSTM/Transformer) trained in Yedetore et al. (2023) with different random seeds, we use the model with the best test perplexity.

[11]We added this Transformer since we wanted to evaluate the information in the English Wikipedia training corpus (the most realistic developmentally in terms of size of all the training corpora under consideration) using a more current architecture than the LSTM that WFL use. We used one of the large Transformer architectures used in Yedetore et al., 2023: 8 layers, hidden and embedding size 1600, and 16 attention heads, trained using the same training regime. Since the current task is limited to single sentences, we lowered the Transformer's context size to 30 (compared to 500 in Yedetore et al. 2023), closer to the average sentence size in the Wikipedia dataset (27.2).

[12]Model version 'text-davinci-003', the latest supported version not fine-tuned using reinforcement learning from human feedback (RLHF) for chat and other applications; however, the model is still trained with supervised fine-tuning, and it is proprietary. See `https://archive.today/2023.10. 07-060351/https://platform.openai.com/docs/models/gpt-3-5` for OpenAI's documentation retrieved October 2023 (archived snapshot).

[13]WFL also use another LSTM, from Jozefowicz et al. (2016). We chose not to include that model in our evaluation since it is extremely slow to work with. For WFL's evaluation, which used a small number of sentences, this was not a problem, but our evaluation relied on a much larger number of sentences, making Jozefowicz et al. (2016)'s model impracticable.

Model	∼Tokens in training data	∼Human equivalent
CHILDES LSTM (Yedetore et al., 2023)	8.6 million	10 months
CHILDES Transformer (Yedetore et al., 2023)		
Wikipedia LSTM (Gulordava et al., 2018)	90 million	8 years
Wikipedia Transformer		
GPT-2 (Radford et al., 2019)	8 billion	730 years
GPT-3 (Brown et al., 2020)	114 billion	10,300 years
GPT-j (Wang and Komatsuzaki, 2021)	402 billion	36,540 years

Table 1: Training data size of the seven language models considered here, and the human linguistic experience equivalent to these data sizes; human equivalents follow WFL (based on Hart and Risley 1992) who assume a daily exposure to ∼30,000 words by children, or around 11 million words per year.

ceptability are aligned, should involve not just a higher probability to the grammatical continuation but a *much* higher one. However, in order to give the models a better chance of refuting the APS, we will adopt a very lenient criterion for success and only ask if the probability assigned to the grammatical continuation is higher than that assigned to the ungrammatical one, without taking into account how much higher it is. This will allow a network to be considered successful even if it prefers the grammatical continuation by the slightest of margins. This lenient condition for success will strengthen our conclusions from cases of failure, which we get to in the sections below: if a network fails even with this lenient condition of success, this failure can be taken seriously.

Here and below we will follow WFL (and the psycholinguistic literature that they build on) and illustrate using *surprisal* values, where the surprisal of x is $-\log P(x)$, which is simply the negative of the logarithmically-scaled probability of x.[14] The lower the probability the higher the surprisal; when the probability approaches 0 the surprisal tends to infinity, and as the probability approaches 1 the surprisal tends to 0. Since higher probability corresponds to lower surprisal, support for the model will come from its assigning lower surprisal to a grammatical continuation than to an ungrammatical one, which, as mentioned, is what WFL indeed find in simple cases.

Figure 1 illustrates the preference of the models considered here for the grammatical continuation over the ungrammatical one in a very simple case by plotting surprisal values for sentences (1a) and (1b). All models assign a lower surprisal value (i.e., a higher proba-

[14]WFL's methodology includes looking not just at $+filler$ cases, as in (1a) and (1b), but also at the corresponding $-filler$ ones, as in (1c) and (1d). We will follow WFL in this in our discussion in sections 5.2 and 5.3 below, but for the present we will attempt to keep the presentation simple by considering only $+filler$ pairs.
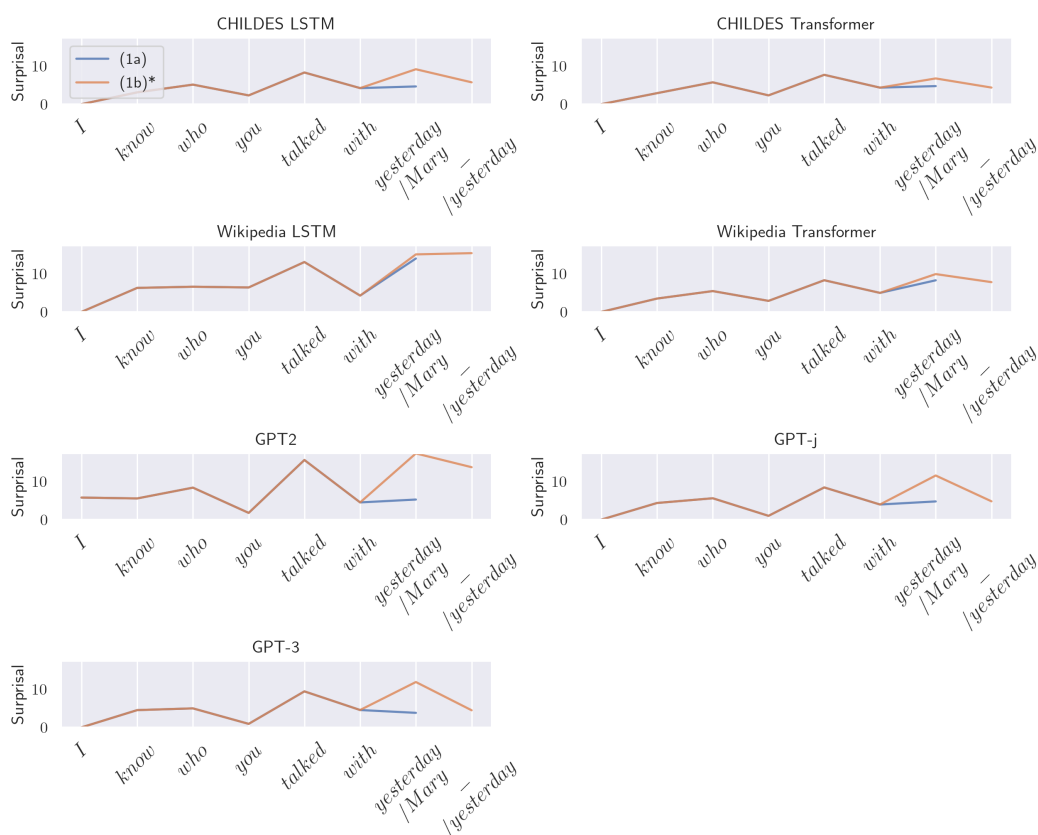
Figure 1: Raw surprisal values outputted by the LLMs for the grammatical (1a), in blue, and ungrammatical (1b), in orange. All models correctly output lower surprisal values for the grammatical continuation.

bility) to the grammatical continuation 'yesterday' in the gapped sentence than to 'Mary'. This suggests that the input is sufficiently rich for a general-purpose learner to acquire from it an approximation of some basic aspects of wh-movement. (We remind the reader that ANNs are used as proxies in the current paper. Their success is not meant to imply that they have acquired the knowledge under consideration and is only used to suggest that there is sufficient information in the data to support the acquisition of this knowledge by a future good model.)

WFL further suggest that the LLMs go beyond the basic knowledge that fillers and gaps go hand in hand. Specifically, they claim that LLMs are aware of *islands* (Ross, 1967): configurations in which a gap is bad even if there is a filler upstream. We illustrate this with the following:

(2)   * I know who [[the question whether __ jumped] surprised Mary yesterday].

While, as discussed above, a filler upstream generally increases the LLMs' expectation of a gap downstream, this expectation should be reduced within the subject of the embed-
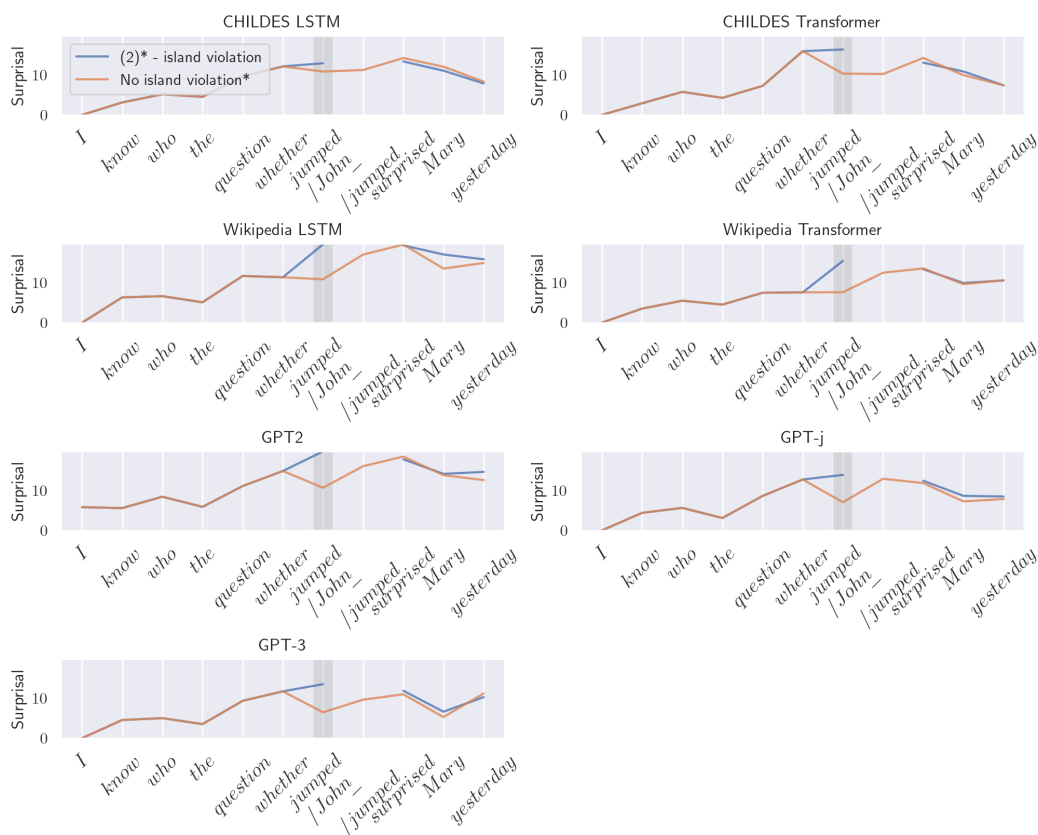
Figure 2: Raw surprisal values for the island violation sentence in (2), in blue, and a variant of the sentence with no island violation (we use 'John' instead of the island-internal gap), in orange. All models are correctly surprised to find a gap within the island. Note that since the variant with 'John' has no downstream gap that would correspond to the upstream filler, it is ungrammatical. For a grammatical version one could replace 'Mary' with a gap. This matter, however, is orthogonal to the surprisal at the island-internal gap, which is what this figure illustrates.

ded clause in (2). This subject is an island to movement, and extraction from within it is unacceptable and presumably highly unlikely. Figure 2 shows that the models are indeed surprised by the gap in (2), suggesting that their training corpora are informative with respect to this aspect of wh-movement.[15]

WFL consider a range of similar cases and conclude that the LLMs are aware of islands quite generally. From this they conclude that linguistically-neutral learners can acquire the intricacies of wh-movement from the input data. In other words, the input data are not

---

[15]The literature discusses various cases in which extraction from subjects (and other islands) is judged acceptable by speakers. Here and below we focus on relatively simple examples in which speaker judgments are clear, and our evaluation will concern the extent to which LLM preferences approximate these clear speaker judgments.

10

impoverished after all with respect to wh-movement, and an APS in this domain falls apart.

We wish to probe the performance of the LLMs — and through that, the richness of the stimulus — further. We will argue that the LLMs do not, in fact, show adequate awareness of wh-movement and therefore do not debunk the APS in this domain. More tentatively, we will provide evidence suggesting that the LLMs' failure is attributable to the input data. If correct, the APS in this domain stands.

4   LLMs fail on slightly more complex (but still simple) cases of wh-movement

We start with a well-studied nuance of islands: in various cases, an otherwise impossible gap inside an island is made possible by a separate gap elsewhere. For example, while (3a), with a subject-internal gap, is bad, its counterpart in (3b), which has an added gap in the direct-object position of the main clause, is good. This phenomenon is known as a *parasitic gap* (PG): the gap inside the subject island becomes acceptable parasitically, based on the direct-object gap.[16]

(3)   a.   * I know who [John's talking to __] is going to annoy you soon.
      b.   I know who [John's talking to __] is going to annoy __ soon.

Somewhat similarly, while (4a), with a gap inside a conjunct, is bad, its counterpart in (4b), where there is a gap in the other conjunct as well, is good. This phenomenon is known as *across-the-board movement* (ATB).[17]

(4)   a.   * I know who John [met __ recently] and [is going to annoy you soon].
      b.   I know who John [met __ recently] and [is going to annoy __ soon].

Do LLMs approximate the patterns of PG and ATB? Figures 3-4 illustrate that all the LLMs under consideration fail on (3a) and (4a), even on our very lenient condition of success: they do not just fail to assign a much higher probability to the grammatical continuation over the ungrammatical one in this simple case; they actually prefer the *ungrammatical* continuation. This seems to indicate that the ANNs have failed to acquire a good approximation of the relevant constructions. At the very least, then, our results challenge WFL's claim that LLMs refute the APS in the domain of wh-movement. More tentatively, the failure of the models suggests that the input is not sufficiently rich to allow a linguistically-neutral learner to acquire wh-movement from it. To the extent that this conclusion holds, we now have a stronger APS than earlier attempts in this domain, since the current version

---

[16]Not all impossible gaps can be rescued in this way. For example, adding further gaps does little to improve (2) above.

[17]We set aside the important question of what stands behind PGs and ATB and whether the two are related. See Ross (1967), Engdahl (1983), Haïk (1985), Williams (1990), Munn (1992), Postal (1993), Nissenbaum (2000), and Hornstein and Nunes (2002), among others, for discussion.
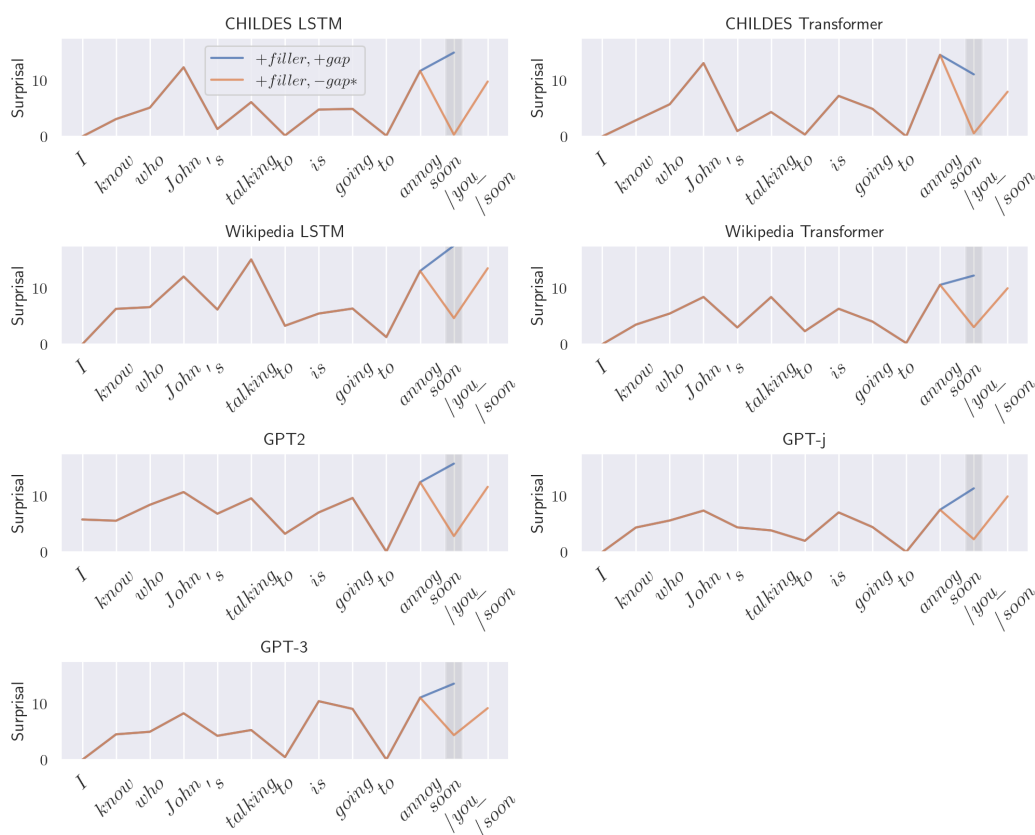
Figure 3: Raw surprisal values for the ungrammatical sentence (3a) which violates a subject island, in orange, and its grammatical variant (3b), in blue, where a parasitic gap makes it possible to escape the island. For measuring the model's expectation for a gap, surprisal is measured at the adverb 'soon', which indicates a gap. This is compared with surprisal at 'John' which plugs the gap at the same position. All networks wrongly assign a higher surprisal value to the grammatical continuation.

is based on general-purpose learners most of which are trained on at least as much linguistic input as children see. In order to support this conclusion, however, we wish to go beyond the examples just discussed and provide evidence that the failure is not due to accidents of the present setup.

5 Failed attempts to rescue the richness of the stimulus in the face of LLM performance

Above we followed WFL in using LLMs to ask whether a linguistically-neutral learner can acquire knowledge of wh-movement from a training corpus that is comparable in size to what children are exposed to. Both for PG and for ATB, and across a broad range of ANNs

Figure 4: Raw surprisal values for the ungrammatical sentence (4a) which violates the coordinate structure constraint (orange), and its grammatical variant (4b) where ATB movement makes it possible to avoid the constraint (blue). All networks wrongly assign a higher surprisal value to the grammatical continuation 'soon' rather than to 'John'.

and training corpora, all the models we looked at preferred the ungrammatical continuation over the corresponding grammatical one. This weakened WFL's claim that LLMs undo the APS in this domain and in fact provided some preliminary support for this APS. If our empirical basis is the previous section alone, however, our conclusions will remain weak. This is so for the following reason: while the failure of a good linguistically-neutral learner on the examples above would indeed be indicative of an impoverished input, it is possible that current ANNs are simply not sufficiently good learners in this regard. For example, they might be sensitive to lexical choices in a way that makes them fail on the specific example we examined above, or they might be biased against the relevant dependencies or possibly even too weak altogether to adequately approximate them. If that happens to be the case, then the LLMs' failure to prefer the grammatical continuation will not teach us about the richness of the input data and will therefore be unhelpful for assessing the APS.

In the present section we attempt to address the general concern about the adequacy of the ANNs, which we break down into three separate investigations. We first ask whether

the failure that we just saw is an accident of the particular lexical choices that we used (section 5.1). We then ask, building on WFL's methodology, whether the failure was due to a general preference for ungapped continuations that is so strong as to override a preference for the correct form (section 5.2). Such a state of affairs could in principle weaken our argument: an ANN might be aware of the correct pattern, but this awareness would be hidden by the preference for ungapped continuations. Finally, we ask whether current LLMs are, for whatever reasons, incapable of approximating the correct preference for the sentences under considerations, even if the input is clearly rich (section 5.3). Each of the three investigations could in principle lead to the conclusion that the ANNs are too weak, and that consequently their failure does not provide significant support to the APS. As we will show, however, none of the investigations actually lead to this conclusion. The failure is not tied to specific accidents of lexical choice (section 5.1). The models fail even when a potential preference for non-gapped continuations is taken into consideration (section 5.2). And at least one model improves significantly once retrained on an enriched corpus, suggesting that the failure on the original corpus reflects the poverty of that corpus (section 5.3). While the opacity of current LLMs requires caution in interpreting their performance, our investigations will strengthen the case that it is not (just) the inadequacies of the models that prevent them from approximating the patterns under consideration and that these patterns are not sufficiently well represented in corpora of the relevant size to justify their acquisition by a linguistically-neutral learner. We will therefore conclude that current LLMs support a qualified APS in the domain of wh-movement.

5.1 Lexical accident?

Our illustration above of how the LLMs prefer the ungrammatical continuation over the grammatical one for ATB and PG used one pair of sentences for each of the two patterns. This raises the obvious worry that the failure of the LLMs reflects some accidents of the specific sentences that we used. This worry is lessened to some extent by the fact that we looked at a broad range of different models trained on different corpora. We find it unlikely that all these models and all these training corpora just happen to have the same blind spot when it comes to the specific sentences that we used above and that otherwise the models approximate the patterns well. Still, it is clearly useful to examine more systematically what happens when we vary the lexical choices for the two patterns.

In order to test the performance of the networks on PG and ATB sentences more broadly, we systematically varied the lexical choices in (3) and (4), repeated here.

(5) a. * I know who [John's talking to __] is going to annoy you soon.

b. I know who [John's talking to __] is going to annoy __ soon.

(6) a. * I know who John [met __ recently] and [is going to annoy you soon].

b. I know who John [met __ recently] and [is going to annoy __ soon].

We generated the sentences by template, using simple context-free grammars. Excerpts from these grammars and a sample of the generated sentences are given in Tables 2 and

Figure 5: Model accuracy on the $\Delta + filler$ condition for the PG and ATB datasets. Accuracy is measured as the ratio of cases where the model assigns a higher probability to the grammatical sentence continuation.

3. The full grammars are given in Appendix A. 8,064 sentence tuples were generated for PG and 6,624 for ATB. For a given model and a given pair of sentences, we looked at the surprisal of the model at the critical point on each member of the pair. For (5), for example, we checked whether after the shared prefix "I know who [John's talking to __] is going to annoy ..." surprisal was higher at the ungapped, ungrammatical continuation 'you' as in (5a) than in the gapped, grammatical continuation 'soon' as in (5b). If it was, and in line with our lenient condition for success that is satisfied by any kind of preference for the grammatical continuation regardless of magnitude, this counted as a success. We will write $\Delta = Surprisal$(ungapped continuation|shared prefix)$-Surprisal$(gapped continuation|shared prefix), and we will write $\Delta_{+filler}$ to indicate that the shared prefix has an upstream filler. Using this notation, we can write the condition for success as $\Delta_{+filler} > 0$.

15

Figure 5 plots the results of examining $\Delta + filler$ preferences for the PG and ATB datasets. In both cases, the best performance by a large margin is that of GPT-3, with 40.9% accuracy on the PG dataset and 71.6% accuracy on the ATB dataset. We are not sure to what extent these numbers can be taken to indicate an approximation of the relevant patterns. If it is a success then it is an underwhelming one. And it is an uninformative one: recall that GPT-3 has been trained on the equivalent of 10,000 years of linguistic experience (and is also further improved manually in various ways), so even if it approximates the relevant patterns, this does not indicate that a general-purpose learner would acquire the relevant knowledge from a developmentally-realistic corpus of just a few years of linguistic experience. Setting GPT-3 aside, the models perform very poorly, with the best performance on PG being Wikipedia LSTM's 18.1% accuracy and the best performance on ATB being CHILDES Transformer's 30.1% accuracy. In other words, the models do not just fail to prefer the grammatical continuation over the ungrammatical one, they positively prefer the ungrammatical continuation in the vast majority of the pairs. Recall WFL's contention that LLMs show that a linguistically-neutral learner can acquire knowledge of wh-movement from a realistic corpus. In the face of our results, WFL's claim needs to be abandoned: current LLMs provide no basis for such a conclusion.

5.2 A preference for ungapped continuations?

In the absence of a single linguistically-neutral learner at present that acquires a reasonable approximation of wh-movement from a realistic corpus, the APS for wh-movement is not challenged. As mentioned, however, we want to go further than this and actually strengthen the APS by showing that not all of the failure of the networks can be attributed to their own failings; some of the blame resides with the lack of sufficient information in the corpus.

Our second investigation, building on WFL's methodology, aims at this strengthening of the APS by focusing on the possibility that the ANNs might be letting irrelevant factors obscure the extent to which they have acquired an approximation of the pattern of wh-movement. Specifically, we will ask whether the networks have a local preference for or against gapped continuations that might make them succeed or fail for the wrong reasons.

Consider again (5) (="I know who [John's talking to __] is going to annoy *you/✓__ soon"). A sufficiently strong local preference about the critical area can affect a given ANN's success regardless of whether it has acquired any approximation of PG, or of wh-movement in general. It could be, for example, that the ANN correctly assigns a higher probability to 'soon' than to 'you' but that it does so because it ignores the filler altogether and simply prefers 'annoy soon' to 'annoy you'. The network's success, then, should not count as evidence for its grasp of wh-movement (and against the APS). Conversely, it could be that the ANN has, in fact, acquired a good knowledge of wh-movement but that it incorrectly prefers 'you' to 'soon' because of similarly irrelevant reasons. For example, perhaps it has a strong preference for ungapped continuations in general, or perhaps it has such a preference in the present case because the lexical frequency of 'you' is very high. The network's failure, then, should not count as evidence against its knowledge of

wh-movement (and for the APS).

To what extent might such local preferences affect the ANNs? We are not entirely sure. A good enough learner would presumably not get confused by such irrelevant factors as lexical frequency and preferences for or against gapped continuations, and the fact that all our models perform well on very simple filler-gap dependencies as in (1) is at least suggestive of their ability to overcome any such confusion when the training data are sufficiently rich. However, beyond this suggestive evidence it is hard to tell whether current ANNs are good enough learners in this sense, and it strikes us as reasonable to further investigate possible confusion by irrelevant factors that might override the preference for the correct pattern.

We will explore the possible effect of irrelevant factors of the kind just mentioned by looking at each LLM's preference for gapped over ungapped continuations and comparing this preference when there is an upstream filler and when there is no such filler. When an upstream filler is present, the model's preference for a gapped continuation (e.g., 'annoy soon') over the ungapped continuation ('annoy you') should be stronger than when an upstream filler is absent. In other words, we will be looking at whole paradigms of the shape we already saw in (1) and not just at those portions of the paradigm in which a filler is present. Such a paradigm is illustrated for PG in (7) and for ATB in (8). Underlined words indicate the ±filler alternations. Words in bold indicate the critical region that shows whether the continuation is gapped or not.

(7) PG

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's talking to is going to annoy **soon**.	*I know <u>who</u> John's talking to is going to annoy **you** soon.
$-filler$	*I know <u>that</u> John's talking to <u>Mary</u> is going to annoy **soon**.	I know <u>that</u> John's talking to <u>Mary</u> is going to annoy **you** soon.

(8) ATB

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John met recently and is going to annoy **soon**.	*I know <u>who</u> John met recently and is going to annoy **you** soon.
$-filler$	*I know <u>that</u> John met <u>Bob</u> recently and is going to annoy **soon**.	I know <u>that</u> John met <u>Bob</u> recently and is going to annoy **you** soon.

Extending our lenient condition for success used above, we will now consider it a success for a given model on a particular paradigm if its preference for the gapped continuation (regardless of its absolute magnitude or even its sign) is higher in the presence of an upstream filler than in its absence. Above we introduced the notation $\Delta = Surprisal(\text{ungapped continuation}|\text{shared prefix}) - Surprisal(\text{gapped continuation}|\text{shared prefix})$ for the extent of the preference for the gapped continuation over the ungapped continuation, and we wrote $\Delta_{+filler}$ when the shared prefix had an upstream filler. We will

now consider also the analogous $\Delta_{-filler}$, for the part of the paradigm where the shared prefix does not have an upstairs filler. And we will consider it a success for the model if $\Delta_{+filler} > \Delta_{-filler}$. This lenient condition of cross-paradigm success follows the logic of difference-in-differences and is very much in line with WFL's evaluation.[18]

In order to test the models across a large number of paradigms, with many different lexical choices, we used the same grammars mentioned in the previous section. In our earlier discussion we used the $+filler$ pairs generated by the grammar. In the present section we use also the corresponding $-filler$ pairs, and from each paradigm of $+filler$ and $-filler$ pairs we compute $\Delta_{\pm filler}$ values. Excerpts from the grammars are provided in Table 2 (for PGs) and Table 3 (for ATB).

Figure 6 plots the LLMs' performance for the cross-paradigm (difference-in-differences) condition. All models except CHILDES LSTM have higher scores for the present measure of $\Delta_{+filler} > \Delta_{-filler}$ than they did for the earlier measure of $\Delta_{+filler} > 0$ (Figure 5), and this holds both for PG and for ATB. This improvement, however, does nothing to endanger the APS. For such a threat to be credible, one would presumably need some kind of linguistically-neutral learner that does at least moderately well in terms of raw ($\Delta_{+filler} > 0$) scores, and so far none of the current models comes close. Even setting that aside and focusing only on the current $\Delta_{+filler} > \Delta_{-filler}$ evaluation, however, we can see that only the extremely large GPT-j and GPT-3 obtain scores that are convincingly high. However, GPT-j is trained on the equivalent of 500 lifetimes of human linguistic exposure. GPT-3 is trained on the equivalent of 100 lifetimes, and fine-tuned further on downstream language tasks. Even GPT-2, trained on the equivalent of ten lifetimes — and thus two orders of magnitude at least above what a child hears by the time they have knowledge of PG and ATB — only reaches modest scores, below 80%. And the smaller models obtain much lower scores. This includes the English Wikipedia LSTM and Transformer, whose training corpus corresponds to about 8 years of linguistic exposure, arguably the most realistic developmentally in terms of size of all the models.[19]

In the present section we tried to give the ANNs — and, through them, the richness of the primary linguistic data — another chance by considering the possibility that the input data are rich but that the ANNs have failed earlier because of a local preference for ungapped continuations. Our criterion for success, $\Delta_{+filler} > \Delta_{-filler}$, only asked that the preference for the gapped continuation over the ungapped continuation be somewhat higher in the presence of an upstream filler. But even with this very weak notion of success, the performance of the networks remained underwhelming.

---

[18]Of course, this new criterion still allows for various irrelevant factors to affect success. For example, a model could become successful simply by deciding that 'who' corresponds to a high probability for 'soon' and a low probability for 'you' anywhere in the sentence and that 'that' corresponds to the opposite. We set aside such worries here.

[19]The gradual improvement of LLM scores as the corpora become very large suggests that current models are in principle capable of improving their approximation of the pattern of wh-movement, but also that this improvement requires much more information than what is present in a corpus that corresponds to anything a child might encounter. We return to this point in the next subsection.

$S \rightarrow \langle PREAMBLE \rangle \; \langle \pm F \rangle \; \langle \pm G \rangle$
$\langle PREAMBLE \rangle \rightarrow$ *I know*
$\langle +F \rangle \rightarrow$ <u>*who*</u> $\langle NAME1 \rangle \langle GEN \rangle \; \langle NP \rangle$
$\langle -F \rangle \rightarrow$ <u>*that*</u> $\langle NAME1 \rangle \langle GEN \rangle \; \langle NP \rangle \; \langle \underline{NAME2} \rangle$
$\langle +G \rangle \rightarrow \langle CONN \rangle \; \langle V \rangle \; \langle \mathbf{ADV} \rangle$
$\langle -G \rangle \rightarrow \langle CONN \rangle \; \langle V \rangle \; \langle \mathbf{OBJ} \rangle \; \langle ADV \rangle$
$\langle GEN \rangle \rightarrow$ *'s*
$\langle NP \rangle \rightarrow \langle NP\_SIMPLE \rangle \; | \; \langle NP\_COMPLEX \rangle$
$\langle NP\_SIMPLE \rangle \rightarrow \langle GERUND \rangle$
$\langle NP\_COMPLEX \rangle \rightarrow \langle N\_EMBEDDED \rangle$ 'to' $\langle V\_EMBEDDED \rangle$
$\langle CONN \rangle \rightarrow$ 'is about to' | 'is likely to' | 'is going to' | 'is expected to'
$\langle V \rangle \rightarrow$ 'bother' | 'annoy' | 'disturb'
$\langle OBJ \rangle \rightarrow$ 'you' | 'us' | 'Kim'
$\langle GERUND \rangle \rightarrow$ 'talking to' | 'dancing with' | 'playing with'
$\langle N\_EMBEDDED \rangle \rightarrow$ 'decision' | 'intent' | 'effort' | 'attempt' | 'failure'
$\langle V\_EMBEDDED \rangle \rightarrow$ 'talk to' | 'call' | 'meet' | 'dance with' | 'play with'
$\langle ADV \rangle \rightarrow$ 'soon' | 'eventually'
. . .
$\Rightarrow$ I know <u>who</u> John's talking to is going to annoy **soon**. $_{(+filler,+gap)}$
$\Rightarrow$ * I know <u>who</u> John's talking to is going to annoy **you** soon. $_{(+filler,-gap)}$
$\Rightarrow$ * I know <u>that</u> John's talking to <u>Mary</u> is going to annoy **soon**. $_{(+filler,-gap)}$
$\Rightarrow$ I know <u>that</u> John's talking to <u>Mary</u> is going to annoy **you** soon. $_{(-filler,-gap)}$

Table 2: Excerpt from the context-free grammar used to generate PG sentences for the experiments in Section 5, and sample sentences generated from it. Underlined words alternate according to the $\pm filler$ condition; words in bold mark the position where the $\pm gap$ condition becomes evident and surprisal is measured.

5.3 A general inability to acquire a suitable preference?

Our third investigation addresses the worry that the ANNs might have failed not because the input was poor but because the networks were too weak to acquire an approximation of the pattern.[20] As with our earlier investigations, we will provide evidence against this possibility, and this evidence, in turn, will significantly strengthen the APS for wh-movement.

In order to make the point that the failure of the ANNs is (also) due to the input and not

---

[20]Recall that throughout this paper we do not ask whether the ANNs have actually acquired the pattern. We are only asking whether they have acquired a passable approximation that, in turn, could suggest that an adequate general-purpose learner would actually acquire the pattern itself. That is, we use ANNs only as a tool to assess how well the pattern is represented in the data.

$S \rightarrow \langle PREAMBLE \rangle \; \langle \pm F \rangle \; \langle CONN \rangle \; \langle \pm G \rangle$
$\langle PREAMBLE \rangle \rightarrow$ *I know*
$\langle +F \rangle \; \rightarrow$ <u>*who*</u> $\langle NAME1 \rangle \; \langle VP1 \rangle \; \langle ADV1 \rangle$
$\langle -F \rangle \; \rightarrow$ <u>*that*</u> $\langle NAME1 \rangle \; \langle VP1 \rangle \; \langle \underline{NAME2} \rangle \; \langle ADV1 \rangle$
$\langle +G \rangle \; \rightarrow \langle CONN \rangle \; \langle VP2 \rangle \; \langle \mathbf{ADV2} \rangle$
$\langle -G \rangle \; \rightarrow \langle CONN \rangle \; \langle VP2 \rangle \; \langle \mathbf{OBJ} \rangle \; \langle ADV2 \rangle$
$\langle CONN \rangle \; \rightarrow$ 'and is going to'
$\langle ADV1 \rangle \; \rightarrow$ 'recently' | 'lately'
$\langle ADV2 \rangle \; \rightarrow$ 'soon' | 'today' | 'now'
$\langle CONN \rangle \; \rightarrow$ 'and is going to'
$\langle VP1 \rangle \; \rightarrow \langle VP1\_SIMPLE \rangle \; | \langle VP1\_COMPLEX \rangle$
$\langle VP1\_SIMPLE \rangle \; \rightarrow$ 'met' | 'saw'
$\langle VP2 \rangle \; \rightarrow \langle VP2\_COMPLEX \rangle \; | \langle VP2\_SIMPLE \rangle$
$\langle VP2\_SIMPLE \rangle \; \rightarrow$ 'hug' | 'slap' | 'kiss'
$\langle OBJ \rangle \; \rightarrow$ 'you' | 'us' | 'Kim'
. . .
$\Rightarrow$ I know <u>who</u> John met recently and is going to hug **soon**. $(+filler,+gap)$
$\Rightarrow$ *I know <u>who</u> John met recently and is going to hug **you** soon. $(+filler,-gap)$
$\Rightarrow$ *I know <u>that</u> John met <u>Bob</u> recently and is going to hug **soon**. $(-filler,+gap)$
$\Rightarrow$ I know <u>that</u> John met <u>Bob</u> recently and is going to hug **you** soon. $(-filler,-gap)$

Table 3: Excerpt from the context-free grammar used to generate ATB sentences for the experiments in Section 4, and sample sentences generated from it. The full grammars are given in Appendix A.

(just) to the ANNs themselves, we checked what happens when the corpus is clearly not impoverished. To do so, we took one of the learners, the Transformer trained on English Wikipedia, and retrained it on a corpus that was identical to its original training dataset, but with an addition of instances of PG and ATB.

These additional instances are generated by template, using a variant of the CFGs that we used in sections 5.1 and 5.2. To increase the probability that an improved performance by the model will reflect generalization rather than memorization, the structure of the additional instances is different from those of the test sentences from Section 5.2. Example training and test sentences are given in Table 4. The full CFGs used in creating the additional instances are given in Appendix C.

From each CFG of each phenomenon (PG/ATB), we sampled 100 sentences for the two grammatical conditions ($+filler, +gap$ and $-filler, -gap$), totaling 200 extra sentences. These sentences were added to the original dataset, and the model was trained using the
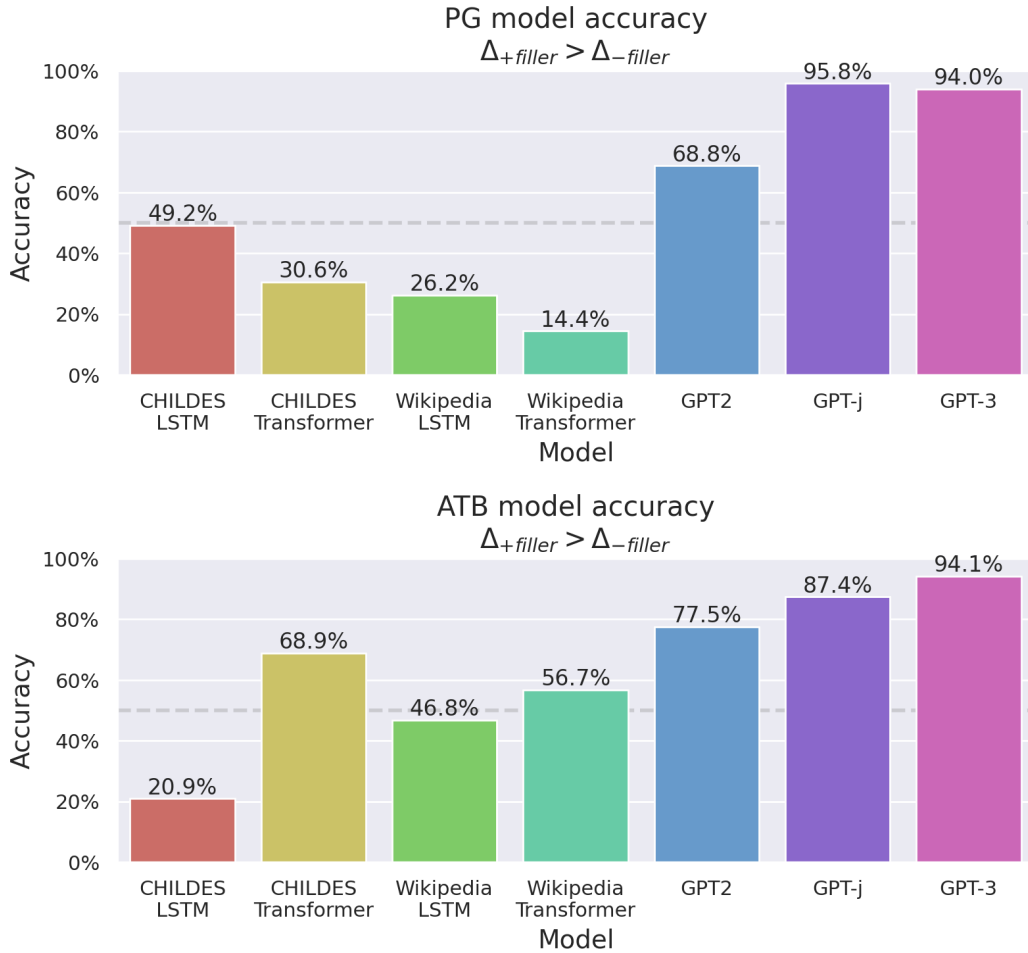
Figure 6: Model accuracy on the difference-in-differences condition for the PG and ATB datasets. Accuracy is measured as the ratio of cases where $\Delta + filler > \Delta - filler$, i.e., when the model shows a relative higher preference for a gap when the gap follows a filler than when it does not.

same regime as in Yedetore et al. (2023) (itself based on the training regime in Gulordava et al., 2018). The model was trained for 48 hours or until reaching the early-stop condition from Yedetore et al. (2023) which stops the training if the validation loss does not improve for more than two consecutive epochs. Due to the long training times of the model, the results reported here are for one random seed with no hyper-parameters search. Since the goal of this experiment was to demonstrate the model's ability to improve significantly given more data, this was sufficient.

The model's performance on the training and test set, before and after retraining, is visualized in Figure 7.

For both ATB and PG the performance of the model improves significantly. For ATB, the raw accuracy score improves from 13.7% to 35.8%, and the difference-in-differences
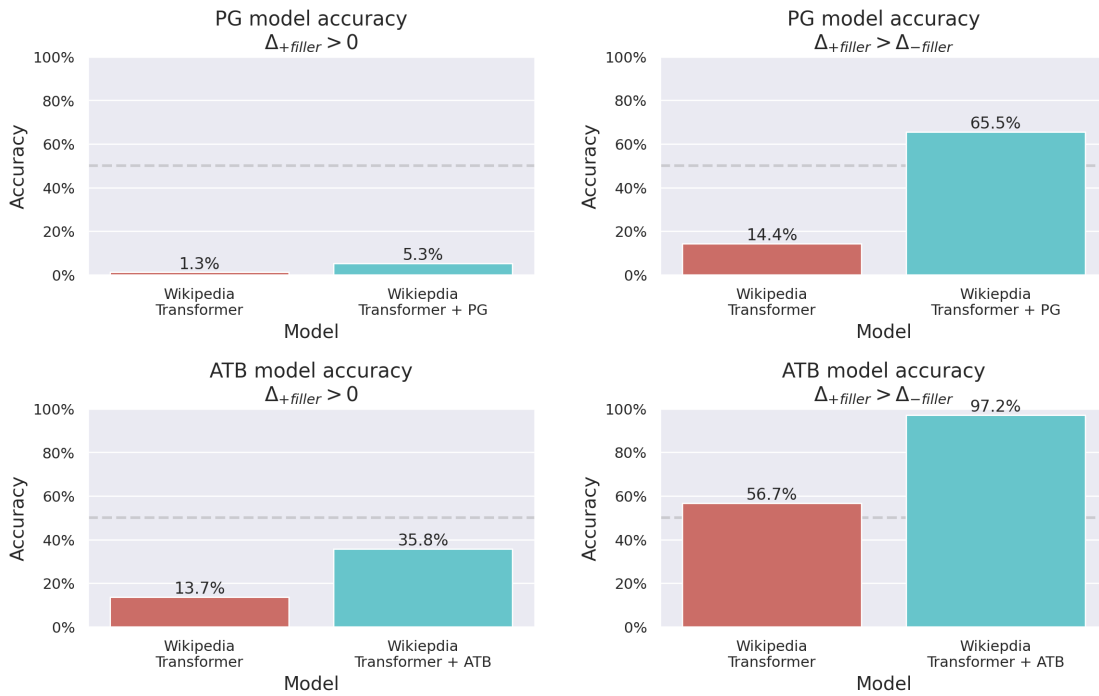
Figure 7: Accuracy for the retrained Transformer model, when trained on the original Wikipedia vs. when trained on the same dataset with extra PG and ATB sentences. The left figures plot accuracy for the $+filler$ condition, and the right figures plot accuracy for the difference-in-difference condition.

PG	**Training Examples** • I know who John's attitude towards upset yesterday. • I know who John's friendship with will annoy soon. • I know who John's praising of amused lately. **Test Examples** • I know who John's talking to is about to bother soon. • I know who John's playing with is going to annoy eventually. • I know who John's failure to dance with is going to disturb soon.
ATB	**Training Examples** • I know who John saw yesterday and kissed today. • I know who John helped recently and married today. • I know who John hugged often and will insult soon. **Test Examples** • I know who John met recently and is going to complain to Patricia about soon. • I know who John said that Mary saw lately and is going to be glad to hug now. • I know who John asked Mary about lately and is going to claim that Patricia will hug today.

Table 4: Example training and test sentences for the retraining task in Section 5.3. A sample of simplified training PG and ATB sentences are added to the model's original training data (English Wikipedia), and the model is then tested on the full battery of sentences from Section 5. The full CFGs for the training and test datasets are given in Appendices A and C.

score improves from 56.7% to 97.2%. This is a dramatic improvement over the performance of the model on its original corpus and is higher than the performance of other architectures when trained on a much larger corpus. For PG, the raw score improves modestly, from 1.3% to 5.3%, while the DID score improves from 14.4% to 65.5%. The raw score for PG certainly doesn't inspire confidence that the model has acquired the dependency. Recall, however, that this is not what we were after here. Our question was whether the model is so weak that its poor performance when trained on the original corpus reflects its inability to do better. The retraining results show that the model can do considerably better once the corpus is sufficiently rich.

Caution is required in interpreting this result. Like all current LLMs, our model is opaque, and we are limited in the conclusions that we can draw from it. In particular, while we observe that it has improved when retrained on a corpus that is enriched in a certain way, it is possible that there are other kinds of evidence for the patterns under consideration that a good general-purpose learner would be able to make use of and that our model cannot. What we found is consistent with such evidence being in the original corpus. Our use of retraining data that were structurally different from the test data was aimed at lessening this

worry, since improvement suggests an ability to generalize and not just memorize. This, in turn, increases the plausibility that the model would have been able to generalize from other kinds of relevant examples if the original corpus had been sufficiently rich. But the opacity of the model prevents us from saying more, and our results here must be qualified accordingly.

6   Conclusion

The APS has been central to linguists' reasoning about innateness for a long time. It has always been difficult, however, to estimate just how much information a linguistically-neutral learner might hope to extract from a realistic input. Modern ANNs promise to change this, and WFL provide a helpful first step in taking advantage of this promise. WFL conclude that the stimulus is rich enough when it comes to wh-movement and that this dismantles the APS in this domain. We showed that this conclusion is premature: even when trained on significantly more linguistic data than children receive and when evaluated using a very lenient condition of success, ANNs fail to reach a passable approximation of the pattern of wh-movement. We illustrated this with parasitic gaps and across-the-board movement.

Is it possible that some future linguistically-neutral learner will succeed where the models that we have examined have failed? Of course. As we mentioned, current models are too opaque and too poorly understood to definitively settle the question of whether the APS for wh-movement holds. We note, however, that the architectures we have considered are generally successful in approximating many other aspects of linguistic data.[21] And some of them have been provided with very generous amounts of linguistic input, in some cases several orders of magnitude beyond what children receive. Given that none of the ANNs reached an adequate approximation of the pattern for the relatively simple examples that we have considered — and given that at least one network did seem capable of improving its approximation when retrained on a clearly rich corpus — we find it likelier that the stimulus is simply too poor to warrant the acquisition of the relevant aspects of knowledge from a corpus that is even remotely realistic developmentally by a linguistically-neutral learner. And if that turns out to be the case, adult speakers' knowledge of these aspects would mean that children are innately endowed in ways that are not linguistically neutral.

---

[21]As noted earlier, there is no obvious sense in which these ANNs have been handicapped with respect to wh-movement; if anything, the evolution of ANNs with the aim of better accounting for linguistic data might have given current models an unfair advantage when it comes to linguistic knowledge and might disqualify them from counting as linguistically neutral. By treating these ANNs as linguistically-neutral we are giving these models the benefit of the doubt.

References

Baker, C. L. 1978. *Introduction to generative transformational syntax*. Englewood Cliffs, New Jersey: Prentice-Hall.

Berwick, Robert C. 2018. Revolutionary new ideas appear infrequently. In *Syntactic Structures after 60 years: The impact of the Chomskyan revolution in linguistics*, ed. Norbert Hornstein, Howard Lasnik, Pritty Patel-Grosz, and Charles Yang, volume 60, 177–194. Walter de Gruyter.

Berwick, Robert C., Paul Pietroski, Beracah Yankama, and Noam Chomsky. 2011. Poverty of the stimulus revisited. *Cognitive Science* 35:1207–1242.

Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, and Amanda Askell. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33:1877–1901.

Chomsky, Noam. 1957. *Syntactic structures*. The Hague: Mouton.

Chomsky, Noam. 1965. *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam. 1971. *Problems of knowledge and freedom: The Russell lectures*. Pantheon Books.

Chomsky, Noam. 1975. *Current issues in linguistic theory*. The Hague: Mouton.

Chomsky, Noam. 1980. *Rules and representations*. New York: Columbia University Press.

Dyer, Fred C., and Jeffrey A. Dickinson. 1994. Development of sun compensation by honeybees: How partially experienced bees estimate the sun's course. *Proceedings of the National Academy of Sciences* 91:4471–4474.

Elman, Jeffrey L., Elizabeth Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett. 1996. *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: The MIT Press.

Engdahl, E. 1983. Parasitic gaps. *Linguistics and Philosophy* 6:5–34.

Foraker, Stephani, Terry Regier, Naveen Khetarpal, Amy Perfors, and Joshua Tenenbaum. 2009. Indirect evidence and the poverty of the stimulus: The case of anaphoric one. *Cognitive Science* 33:287–300.

Gordon, Peter. 1985. Level-ordering in lexical development. *Cognition* 21:73–93.

Gulordava, Kristina, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless Green Recurrent Networks Dream Hierarchically. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1195–1205. New Orleans, Louisiana: Association for Computational Linguistics.

Haïk, Isabelle. 1985. The syntax of operators. Doctoral Dissertation, MIT.

Hart, Betty, and Todd R. Risley. 1992. American parenting of language-learning children: Persisting differences in family-child interactions observed in natural home environments. *Developmental Psychology* 28:1096–1105.

Hornstein, Norbert, and Jairo Nunes. 2002. On asymmetries between parasitic gap and across-the-board constructions. *Syntax* 5.

Hsu, Anne S., and Nick Chater. 2010. The logical problem of language acquisition: A probabilistic perspective. *Cognitive Science* 34:972–1016.

Jozefowicz, Rafal, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the Limits of Language Modeling. *arXiv:1602.02410 [cs]* .

Legate, Julie Anne, and Charles Yang. 2002. Empirical re-assessment of stimulus poverty arguments. *The Linguistic Review* 19.

Lewis, John D, and Jeffery L Elman. 2001. A connectionist investigtion of linguistic arguments from the poverty of the stimulus: Learning the unlearnable. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 23.

Lidz, Jeffrey, Sandra Waxman, and Jennifer Freedman. 2003. What infants know about syntax but couldn't have learned: Experimental evidence for syntactic structure at 18 months. *Cognition* 89:B65–B73.

MacWhinney, Brian. 2014. *The Childes Project: Tools for Analyzing Talk, Volume II: The Database*. New York: Psychology Press, 3 edition.

Munn, Alan. 1992. A null operator analysis of ATB gaps. *The Linguistic Review* 9:1–26.

Nissenbaum, Jon. 2000. Investigations of covert phrase movement. Doctoral Dissertation, MIT, Cambridge, Mass.

Pearl, Lisa, and Jon Sprouse. 2013. Syntactic islands and learning biases: Combining experimental syntax and computational modeling to investigate the language acquisition problem. *Language Acquisition* 20:23–68.

Perfors, Amy, Joshua Tenenbaum, and Terry Regier. 2011. The learnability of abstract syntactic principles. *Cognition* 118:306–338.

Phillips, Colin. 2013. On the nature of island constraints II: Language learning and innateness. In *Experimental syntax and island effects*, ed. Jon Sprouse and Norbert Hornstein, 132–157. Cambridge University Press.

Postal, Paul M. 1993. Parasitic gaps and the across-the-board phenomenon. *Linguistic Inquiry* 24:735–754.

Pullum, Geoffrey, and Barbara Scholz. 2002. Empirical assessment of stimulus poverty arguments. *The Linguistic Review* 19:9–50.

Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1:9.

Rawski, Jonathan, and Jeffrey Heinz. 2019. No free lunch in linguistics or machine learning: Response to pater. *Language* 95:e125–e135.

Reali, Florencia, and Morten Christiansen. 2005. Uncovering the richness of the stimulus: Structure dependence and indirect statistical evidence. *Cognitive Science* 29:1007–1028.

Ross, John R. 1967. Constraints on variables in syntax. Doctoral Dissertation, MIT, Cambridge, MA.

Sprouse, Jon, Beracah Yankama, Sagar Indurkhya, Sandiway Fong, and Robert C. Berwick. 2018. Colorless green ideas do sleep furiously: gradient acceptability and the nature of the grammar. *The Linguistic Review* 35:575–599.

Wang, Ben, and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 billion parameter autoregressive language model.

Warstadt, Alex, and Samuel R Bowman. 2022. What artificial neural networks can tell us about human language acquisition. In *Algebraic structures in natural language*, 17–60. CRC Press.

Wilcox, Ethan Gotlieb, Richard Futrell, and Roger Levy. 2022. Using computational models to test syntactic learnability. *Linguistic Inquiry* 1–88.

Williams, Edwin. 1990. The ATB-theory of parasitic gaps. *The Linguistic Review* 6:265–279.

Wilson, Colin. 2006. Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science* 30:945–982.

Yedetore, Aditya, Tal Linzen, Robert Frank, and R. Thomas McCoy. 2023. How poor is the stimulus? Evaluating hierarchical generalization in neural networks trained on child-directed speech.

A Appendix: context-free grammars

A.1 PG

$\langle S \rangle \rightarrow \langle S\_FG \rangle$

$\langle S\_FG \rangle \rightarrow \langle PREAMBLE \rangle \langle F \rangle \langle G \rangle$

$\langle S\_XG \rangle \rightarrow \langle UNGRAMMATICAL \rangle \langle PREAMBLE \rangle \langle XF \rangle \langle G \rangle$

$\langle S\_FX \rangle \rightarrow \langle UNGRAMMATICAL \rangle \langle PREAMBLE \rangle \langle F \rangle \langle XG \rangle$

$\langle S\_XX \rangle \rightarrow \langle PREAMBLE \rangle \langle XF \rangle \langle XG \rangle$

$\langle UNGRAMMATICAL \rangle \rightarrow$ '*'

$\langle GEN \rangle \rightarrow$ ''s'

$\langle OBJ \rangle \rightarrow$ 'you' | 'us' | 'Kim'

$\langle NAME1 \rangle \rightarrow$ 'Bob' | 'John'

$\langle NAME2 \rangle \rightarrow$ 'Mary' | 'Jennifer'

$\langle NAME3 \rangle \rightarrow$ 'James' | 'Michael'

$\langle NAME4 \rangle \rightarrow$ 'Patricia' | 'Linda'

$\langle PREAMBLE \rangle \rightarrow$ 'I know'

$\langle F \rangle \rightarrow$ 'who' $\langle NAME1 \rangle \langle GEN \rangle \langle NP \rangle$

$\langle XF \rangle \rightarrow$ 'that' $\langle NAME1 \rangle \langle GEN \rangle \langle NP \rangle \langle NAME2 \rangle$

$\langle G \rangle \rightarrow \langle CONN \rangle \langle V \rangle \langle ADV2 \rangle$

$\langle XG \rangle \rightarrow \langle CONN \rangle \langle V \rangle \langle OBJ \rangle \langle ADV2 \rangle$

$\langle NP \rangle \rightarrow \langle NP\_SIMPLE \rangle$ | $\langle NP\_COMPLEX \rangle$

$\langle NP\_SIMPLE \rangle \rightarrow \langle GERUND \rangle$

$\langle NP\_COMPLEX \rangle \rightarrow \langle N\_EMBEDDED \rangle$ 'to' $\langle V\_EMBEDDED \rangle$

$\langle CONN \rangle \rightarrow$ 'is about to' | 'is likely to' | 'is going to' | 'is expected to'

$\langle V \rangle \rightarrow$ 'bother' | 'annoy' | 'disturb'

$\langle GERUND \rangle \rightarrow$ 'talking to' | 'dancing with' | 'playing with'

$\langle N\_EMBEDDED \rangle \rightarrow$ 'decision' | 'intent' | 'effort' | 'attempt' | 'failure'

⟨*V_EMBEDDED*⟩ → 'talk to' | 'call' | 'meet' | 'dance with' | 'play with'
⟨*ADV1*⟩ → 'recently' | 'earlier'
⟨*ADV2*⟩ → 'soon' | 'eventually'

A.2 ATB

⟨*S*⟩ → ⟨*S_FG*⟩
⟨*S_FG*⟩ → ⟨*PREAMBLE*⟩ ⟨*F*⟩ ⟨*G*⟩
⟨*S_XG*⟩ → ⟨*UNGRAMMATICAL*⟩ ⟨*PREAMBLE*⟩ ⟨*XF*⟩ ⟨*G*⟩
⟨*S_FX*⟩ → ⟨*UNGRAMMATICAL*⟩ ⟨*PREAMBLE*⟩ ⟨*F*⟩ ⟨*XG*⟩
⟨*S_XX*⟩ → ⟨*PREAMBLE*⟩ ⟨*XF*⟩ ⟨*XG*⟩
⟨*UNGRAMMATICAL*⟩ → '*'
⟨*GEN*⟩ → ''s'
⟨*OBJ*⟩ → 'you' | 'us' | 'Kim'
⟨*NAME1*⟩ → 'John'
⟨*NAME2*⟩ → 'Mary'
⟨*NAME3*⟩ → 'Bob'
⟨*NAME4*⟩ → 'Patricia'
⟨*PREAMBLE*⟩ → 'I know'
⟨*F*⟩ → 'who' ⟨*NAME1*⟩ ⟨*VP1*⟩ ⟨*ADV1*⟩
⟨*XF*⟩ → 'that' ⟨*NAME1*⟩ ⟨*VP1*⟩ ⟨*NAME3*⟩ ⟨*ADV1*⟩
⟨*G*⟩ → ⟨*CONN*⟩ ⟨*VP2*⟩ ⟨*ADV2*⟩
⟨*XG*⟩ → ⟨*CONN*⟩ ⟨*VP2*⟩ ⟨*OBJ*⟩ ⟨*ADV2*⟩
⟨*ADV1*⟩ → 'recently' | 'lately'
⟨*ADV2*⟩ → 'soon' | 'today' | 'now'
⟨*CONN*⟩ → 'and is going to'
⟨*VP1*⟩ → ⟨*VP1_SIMPLE*⟩ | ⟨*VP1_COMPLEX*⟩
⟨*VP1_SIMPLE*⟩ → 'met' | 'saw'
⟨*VP1_COMPLEX*⟩ → ⟨*VP1_ABOUT*⟩ | ⟨*VP1_TO*⟩ | ⟨*VP1_ADJ*⟩ | ⟨*VP1_EMBEDDED*⟩
⟨*VP1_ABOUT*⟩ → ⟨*V_ABOUT_PAST*⟩ ⟨*NAME2*⟩ 'about'
⟨*V_ABOUT_PAST*⟩ → 'asked' | 'told'
⟨*VP1_TO*⟩ → ⟨*V_TO_PAST*⟩ ⟨*NAME2*⟩ 'to' ⟨*V_TRANS_INF_TO*⟩
⟨*V_TO_PAST*⟩ → 'wanted' | 'asked'
⟨*V_TRANS_INF_TO*⟩ → 'call' | 'invite'
⟨*VP1_ADJ*⟩ → 'was' ⟨*ADJ1*⟩ 'to' ⟨*V_TRANS_INF_ADJ*⟩
⟨*ADJ1*⟩ → 'eager' | 'happy'
⟨*V_TRANS_INF_ADJ*⟩ → 'meet' | 'see'
⟨*VP1_EMBEDDED*⟩ → ⟨*V_EMBEDDING_PAST*⟩ 'that' ⟨*NAME2*⟩ ⟨*V_TRANS_PAST_EMBE*...⟩
⟨*V_EMBEDDING_PAST*⟩ → 'said' | 'insisted'
⟨*V_TRANS_PAST_EMBEDDED*⟩ → 'met' | 'saw'
⟨*VP2*⟩ → ⟨*VP2_COMPLEX*⟩ | ⟨*VP2_SIMPLE*⟩

$\langle VP2\_SIMPLE \rangle \rightarrow$ 'hug' | 'slap' | 'kiss'
$\langle VP2\_COMPLEX \rangle \rightarrow \langle VP2\_ABOUT \rangle | \langle VP2\_TO \rangle | \langle VP2\_ADJ \rangle | \langle VP2\_EMBEDDED \rangle$
$\langle VP2\_ABOUT \rangle \rightarrow \langle V\_ABOUT\_FUTURE \rangle$ 'to' $\langle NAME4 \rangle$ 'about'
$\langle V\_ABOUT\_FUTURE \rangle \rightarrow$ 'complain' | 'write'
$\langle VP2\_TO \rangle \rightarrow \langle V\_TO\_FUTURE \rangle \langle NAME4 \rangle$ 'to' $\langle V\_TRANS\_INF\_TO\_FUTURE \rangle$
$\langle V\_TO\_FUTURE \rangle \rightarrow$ 'encourage' | 'beg'
$\langle V\_TRANS\_INF\_TO\_FUTURE \rangle \rightarrow$ 'hug' | 'slap' | 'kiss'
$\langle VP2\_ADJ \rangle \rightarrow$ 'be' $\langle ADJ2 \rangle$ 'to' $\langle V\_TRANS\_INF\_ADJ2 \rangle$
$\langle ADJ2 \rangle \rightarrow$ 'afraid' | 'glad'
$\langle V\_TRANS\_INF\_ADJ2 \rangle \rightarrow$ 'hug' | 'slap' | 'kiss'
$\langle VP2\_EMBEDDED \rangle \rightarrow \langle V\_EMBEDDING\_FUTURE \rangle$ 'that' $\langle NAME4 \rangle$ 'will' $\langle V\_TRANS\_FUT$
$\langle V\_EMBEDDING\_FUTURE \rangle \rightarrow$ 'claim' | 'predict'
$\langle V\_TRANS\_FUTURE \rangle \rightarrow$ 'hug' | 'slap' | 'kiss'

B Appendix: model failures

Worst 5 four-tuples of sentences per phenomenon (PG, ATB), per model (CHILDES LSTM/Transformer, Wikipedia LSTM/Transformer, GPT-2, GPT-j, GPT-3). Surprisal values are given in parentheses at the relevant position.

B.1 PG – CHILDES LSTM

(9)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's effort to call is going to bother **eventually (14.75)**	*I know <u>who</u> Bob's effort to call is going to bother **you (0.64)** eventually
$-filler$	*I know <u>that</u> Bob's effort to call Mary is going to bother **eventually (12.74)**	I know <u>that</u> Bob's effort to call Mary is going to bother **you (1.70)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -3.06$		

(10)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's effort to call is going to annoy **eventually (15.82)**	*I know <u>who</u> Bob's effort to call is going to annoy **you (0.39)** eventually
$-filler$	*I know <u>that</u> Bob's effort to call Mary is going to annoy **eventually (13.65)**	I know <u>that</u> Bob's effort to call Mary is going to annoy **you (1.23)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -3.01$		

(11)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's effort to call is going to bother **soon (13.28)**	*I know <u>who</u> Bob's effort to call is going to bother **you (0.64)** soon
$-filler$	*I know <u>that</u> Bob's effort to call Mary is going to bother **soon (11.53)**	I know <u>that</u> Bob's effort to call Mary is going to bother **you (1.70)** soon
$\Delta_{-filler} - \Delta_{+filler} = -2.80$		

(12)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's dancing with is going to disturb **eventually (15.03)**	*I know <u>who</u> John's dancing with is going to disturb **us (1.82)** eventually
$-filler$	*I know <u>that</u> John's dancing with Mary is going to disturb **eventually (13.69)**	I know <u>that</u> John's dancing with Mary is going to disturb **us (3.27)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -2.80$		

(13)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's failure to meet is going to disturb **eventually (13.66)**	*I know <u>who</u> Bob's failure to meet is going to disturb **us (1.63)** eventually
$-filler$	*I know <u>that</u> Bob's failure to meet Mary is going to disturb **eventually (12.41)**	I know <u>that</u> Bob's failure to meet Mary is going to disturb **us (3.17)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -2.79$		

B.2   PG – CHILDES Transformer

(14)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's attempt to play with is going to bother **soon (12.45)**	*I know <u>who</u> Bob's attempt to play with is going to bother **us (1.55)** soon
$-filler$	*I know <u>that</u> Bob's attempt to play with Jennifer is going to bother **soon (10.69)**	I know <u>that</u> Bob's attempt to play with Jennifer is going to bother **us (3.20)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.41$		

(15)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's effort to play with is going to bother **soon (12.27)**	*I know <u>who</u> Bob's effort to play with is going to bother **us (1.02)** soon
$-filler$	*I know <u>that</u> Bob's effort to play with Jennifer is going to bother **soon (10.31)**	I know <u>that</u> Bob's effort to play with Jennifer is going to bother **us (2.32)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.27$		

(16)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's attempt to play with is about to bother **soon (12.78)**	*I know <u>who</u> Bob's attempt to play with is about to bother **us (1.63)** soon
$-filler$	*I know <u>that</u> Bob's attempt to play with Jennifer is about to bother **soon (10.91)**	I know <u>that</u> Bob's attempt to play with Jennifer is about to bother **us (2.89)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.13$		

(17)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's attempt to play with is expected to bother **soon (13.39)**	*I know <u>who</u> Bob's attempt to play with is expected to bother **us (2.02)** soon
$-filler$	*I know <u>that</u> Bob's attempt to play with Jennifer is expected to bother **soon (11.92)**	I know <u>that</u> Bob's attempt to play with Jennifer is expected to bother **us (3.66)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.10$		

(18)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's failure to play with is going to bother **soon (12.32)**	*I know <u>who</u> Bob's failure to play with is going to bother **us (1.80)** soon
$-filler$	*I know <u>that</u> Bob's failure to play with Jennifer is going to bother **soon (10.63)**	I know <u>that</u> Bob's failure to play with Jennifer is going to bother **us (3.14)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.03$		

B.3 PG – Wikipedia LSTM

(19)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's intent to call is going to bother **eventually (16.59)**	*I know <u>who</u> John's intent to call is going to bother **us (4.90)** eventually
$-filler$	*I know <u>that</u> John's intent to call Jennifer is going to bother **eventually (13.47)**	I know <u>that</u> John's intent to call Jennifer is going to bother **us (7.20)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -5.42$		

(20)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's failure to call is going to bother **eventually (16.07)**	*I know <u>who</u> John's failure to call is going to bother **us (4.93)** eventually
$-filler$	*I know <u>that</u> John's failure to call Jennifer is going to bother **eventually (13.77)**	I know <u>that</u> John's failure to call Jennifer is going to bother **us (7.89)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -5.26$		

(21)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's failure to call is going to bother **soon (15.47)**	*I know <u>who</u> John's failure to call is going to bother **us (4.93)** soon
$-filler$	*I know <u>that</u> John's failure to call Jennifer is going to bother **soon (13.68)**	I know <u>that</u> John's failure to call Jennifer is going to bother **us (7.89)** soon
$\Delta_{-filler} - \Delta_{+filler} = -4.74$		

(22)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's attempt to talk to is going to bother **eventually (15.99)**	*I know <u>who</u> John's attempt to talk to is going to bother **us (4.42)** eventually
$-filler$	*I know <u>that</u> John's attempt to talk to Jennifer is going to bother **eventually (14.78)**	I know <u>that</u> John's attempt to talk to Jennifer is going to bother **us (7.65)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -4.44$		

(23)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John's intent to talk to is going to bother **eventually (16.62)**	*I know <u>who</u> John's intent to talk to is going to bother **you (5.39)** eventually
	$-filler$	*I know <u>that</u> John's intent to talk to Jennifer is going to bother **eventually (14.89)**	I know <u>that</u> John's intent to talk to Jennifer is going to bother **you (8.06)** eventually
	$\Delta_{-filler} - \Delta_{+filler} = -4.41$		

B.4 PG – Wikipedia Transformer

(24)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> Bob's attempt to talk to is about to bother **soon (11.09)**	*I know <u>who</u> Bob's attempt to talk to is about to bother **you (3.73)** soon
	$-filler$	*I know <u>that</u> Bob's attempt to talk to Jennifer is about to bother **soon (7.94)**	I know <u>that</u> Bob's attempt to talk to Jennifer is about to bother **you (10.53)** soon
	$\Delta_{-filler} - \Delta_{+filler} = -9.95$		

(25)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> Bob's attempt to talk to is expected to bother **soon (10.01)**	*I know <u>who</u> Bob's attempt to talk to is expected to bother **you (3.49)** soon
	$-filler$	*I know <u>that</u> Bob's attempt to talk to Jennifer is expected to bother **soon (7.23)**	I know <u>that</u> Bob's attempt to talk to Jennifer is expected to bother **you (10.49)** soon
	$\Delta_{-filler} - \Delta_{+filler} = -9.78$		

(26)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> Bob's attempt to talk to is expected to bother **eventually (10.76)**	*I know <u>who</u> Bob's attempt to talk to is expected to bother **you (3.49)** eventually
	$-filler$	*I know <u>that</u> Bob's attempt to talk to Jennifer is expected to bother **eventually (8.07)**	I know <u>that</u> Bob's attempt to talk to Jennifer is expected to bother **you (10.49)** eventually
	$\Delta_{-filler} - \Delta_{+filler} = -9.70$		

(27)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's intent to talk to is expected to bother **eventually (10.98)**	*I know <u>who</u> John's intent to talk to is expected to bother **us (4.03)** eventually
$-filler$	*I know <u>that</u> John's intent to talk to Jennifer is expected to bother **eventually (8.32)**	I know <u>that</u> John's intent to talk to Jennifer is expected to bother **us (11.06)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -9.69$		

(28)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's decision to talk to is about to bother **soon (10.66)**	*I know <u>who</u> Bob's decision to talk to is about to bother **you (3.70)** soon
$-filler$	*I know <u>that</u> Bob's decision to talk to Jennifer is about to bother **soon (7.83)**	I know <u>that</u> Bob's decision to talk to Jennifer is about to bother **you (10.50)** soon
$\Delta_{-filler} - \Delta_{+filler} = -9.63$		

B.5 PG – GPT-2

(29)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's talking to is going to annoy **eventually (16.51)**	*I know <u>who</u> Bob's talking to is going to annoy **you (2.77)** eventually
$-filler$	*I know <u>that</u> Bob's talking to Jennifer is going to annoy **eventually (16.17)**	I know <u>that</u> Bob's talking to Jennifer is going to annoy **you (4.70)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -2.26$		

(30)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's decision to meet is about to bother **eventually (18.76)**	*I know <u>who</u> Bob's decision to meet is about to bother **you (2.17)** eventually
$-filler$	*I know <u>that</u> Bob's decision to meet Jennifer is about to bother **eventually (18.42)**	I know <u>that</u> Bob's decision to meet Jennifer is about to bother **you (4.04)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -2.20$		

(31)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's decision to call is likely to disturb **soon (14.28)**	*I know <u>who</u> John's decision to call is likely to disturb **you (2.73)** soon
$-filler$	*I know <u>that</u> John's decision to call Mary is likely to disturb **soon (15.19)**	I know <u>that</u> John's decision to call Mary is likely to disturb **you (5.84)** soon
$\Delta_{-filler} - \Delta_{+filler} = -2.19$		

(32)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's talking to is likely to disturb **eventually (16.82)**	*I know <u>who</u> John's talking to is likely to disturb **you (1.81)** eventually
$-filler$	*I know <u>that</u> John's talking to Jennifer is likely to disturb **eventually (16.53)**	I know <u>that</u> John's talking to Jennifer is likely to disturb **you (3.68)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -2.16$		

(33)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's decision to call is likely to disturb **soon (14.31)**	*I know <u>who</u> Bob's decision to call is likely to disturb **you (2.70)** soon
$-filler$	*I know <u>that</u> Bob's decision to call Mary is likely to disturb **soon (15.37)**	I know <u>that</u> Bob's decision to call Mary is likely to disturb **you (5.91)** soon
$\Delta_{-filler} - \Delta_{+filler} = -2.15$		

B.6 PG – GPT-j

(34)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's attempt to talk to is about to annoy **soon (11.48)**	*I know <u>who</u> John's attempt to talk to is about to annoy **you (1.62)** soon
$-filler$	*I know <u>that</u> John's attempt to talk to Mary is about to annoy **soon (12.86)**	I know <u>that</u> John's attempt to talk to Mary is about to annoy **you (4.70)** soon
$\Delta_{-filler} - \Delta_{+filler} = -1.70$		

(35)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's failure to dance with is about to annoy **eventually (13.09)**	*I know <u>who</u> John's failure to dance with is about to annoy **you (1.53)** eventually
$-filler$	*I know <u>that</u> John's failure to dance with Mary is about to annoy **eventually (13.60)**	I know <u>that</u> John's failure to dance with Mary is about to annoy **you (3.69)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -1.66$		

(36)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's attempt to talk to is about to annoy **eventually (12.23)**	*I know <u>who</u> John's attempt to talk to is about to annoy **you (1.62)** eventually
$-filler$	*I know <u>that</u> John's attempt to talk to Mary is about to annoy **eventually (13.71)**	I know <u>that</u> John's attempt to talk to Mary is about to annoy **you (4.70)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -1.59$		

(37)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's decision to dance with is about to annoy **eventually (13.70)**	*I know <u>who</u> John's decision to dance with is about to annoy **you (1.26)** eventually
$-filler$	*I know <u>that</u> John's decision to dance with Mary is about to annoy **eventually (13.85)**	I know <u>that</u> John's decision to dance with Mary is about to annoy **you (2.95)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -1.54$		

(38)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's decision to dance with is about to annoy **soon (12.74)**	*I know <u>who</u> John's decision to dance with is about to annoy **you (1.26)** soon
$-filler$	*I know <u>that</u> John's decision to dance with Mary is about to annoy **soon (12.92)**	I know <u>that</u> John's decision to dance with Mary is about to annoy **you (2.95)** soon
$\Delta_{-filler} - \Delta_{+filler} = -1.51$		

B.7 PG – GPT-3

(39)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's talking to is likely to annoy **soon (14.53)**	*I know <u>who</u> Bob's talking to is likely to annoy **you (3.39)** soon
$-filler$	*I know <u>that</u> Bob's talking to Mary is likely to annoy **soon (13.53)**	I know <u>that</u> Bob's talking to Mary is likely to annoy **you (7.20)** soon
$\Delta_{-filler} - \Delta_{+filler} = -4.81$		

(40)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's talking to is likely to annoy **soon (14.06)**	*I know <u>who</u> John's talking to is likely to annoy **you (3.55)** soon
$-filler$	*I know <u>that</u> John's talking to Mary is likely to annoy **soon (13.59)**	I know <u>that</u> John's talking to Mary is likely to annoy **you (7.60)** soon
$\Delta_{-filler} - \Delta_{+filler} = -4.52$		

(41)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's playing with is likely to annoy **soon (13.36)**	*I know <u>who</u> Bob's playing with is likely to annoy **you (2.75)** soon
$-filler$	*I know <u>that</u> Bob's playing with Mary is likely to annoy **soon (13.46)**	I know <u>that</u> Bob's playing with Mary is likely to annoy **you (6.94)** soon
$\Delta_{-filler} - \Delta_{+filler} = -4.09$		

(42)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John's talking to is expected to annoy **soon (12.97)**	*I know <u>who</u> John's talking to is expected to annoy **you (3.05)** soon
$-filler$	*I know <u>that</u> John's talking to Mary is expected to annoy **soon (12.61)**	I know <u>that</u> John's talking to Mary is expected to annoy **you (6.56)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.87$		

(43)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> Bob's talking to is likely to annoy **eventually (15.21)**	*I know <u>who</u> Bob's talking to is likely to annoy **you (3.39)** eventually
$-filler$	*I know <u>that</u> Bob's talking to Mary is likely to annoy **eventually (15.18)**	I know <u>that</u> Bob's talking to Mary is likely to annoy **you (7.20)** eventually
$\Delta_{-filler} - \Delta_{+filler} = -3.84$		

B.8 ATB – CHILDES LSTM

(44)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about lately and is going to encourage Patricia to slap **now (9.84)**	*I know <u>who</u> John told Mary about lately and is going to encourage Patricia to slap **you (1.88)** now
$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to encourage Patricia to slap **now (9.24)**	I know <u>that</u> John told Mary about Bob lately and is going to encourage Patricia to slap **you (3.35)** now
$\Delta_{-filler} - \Delta_{+filler} = -2.08$		

(45)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John insisted that Mary met recently and is going to be afraid to kiss **now (8.53)**	*I know <u>who</u> John insisted that Mary met recently and is going to be afraid to kiss **Kim (10.20)** now
$-filler$	*I know <u>that</u> John insisted that Mary met Bob recently and is going to be afraid to kiss **now (8.46)**	I know <u>that</u> John insisted that Mary met Bob recently and is going to be afraid to kiss **Kim (12.17)** now
$\Delta_{-filler} - \Delta_{+filler} = -2.04$		

(46)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John insisted that Mary met recently and is going to be afraid to kiss **today (8.56)**	*I know <u>who</u> John insisted that Mary met recently and is going to be afraid to kiss **Kim (10.20)** today
$-filler$	*I know <u>that</u> John insisted that Mary met Bob recently and is going to be afraid to kiss **today (8.53)**	I know <u>that</u> John insisted that Mary met Bob recently and is going to be afraid to kiss **Kim (12.17)** today
$\Delta_{-filler} - \Delta_{+filler} = -2.00$		

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John insisted that Mary met recently and is going to be afraid to slap **now (9.34)**	*I know <u>who</u> John insisted that Mary met recently and is going to be afraid to slap **Kim (9.90)** now
$-filler$	*I know <u>that</u> John insisted that Mary met Bob recently and is going to be afraid to slap **now (9.29)**	I know <u>that</u> John insisted that Mary met Bob recently and is going to be afraid to slap **Kim (11.85)** now
$\Delta_{-filler} - \Delta_{+filler} = -2.00$		

(47)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John wanted Mary to invite recently and is going to be afraid to slap **now (9.41)**	*I know <u>who</u> John wanted Mary to invite recently and is going to be afraid to slap **Kim (10.37)** now
$-filler$	*I know <u>that</u> John wanted Mary to invite Bob recently and is going to be afraid to slap **now (9.22)**	I know <u>that</u> John wanted Mary to invite Bob recently and is going to be afraid to slap **Kim (12.16)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.98$		

(48)

B.9 ATB – CHILDES Transformer

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John asked Mary about lately and is going to be afraid to slap **now (10.21)**	*I know <u>who</u> John asked Mary about lately and is going to be afraid to slap **us (3.41)** now
$-filler$	*I know <u>that</u> John asked Mary about Bob lately and is going to be afraid to slap **now (9.42)**	I know <u>that</u> John asked Mary about Bob lately and is going to be afraid to slap **us (4.14)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.51$		

(49)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John said that Mary saw recently and is going to slap **now (10.72)**	*I know <u>who</u> John said that Mary saw recently and is going to slap **us (3.14)** now
$-filler$	*I know <u>that</u> John said that Mary saw Bob recently and is going to slap **now (10.18)**	I know <u>that</u> John said that Mary saw Bob recently and is going to slap **us (4.00)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.40$		

(50)

(51)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John said that Mary saw lately and is going to be afraid to slap **now (9.88)**	*I know <u>who</u> John said that Mary saw lately and is going to be afraid to slap **us (3.24)** now
$-filler$	*I know <u>that</u> John said that Mary saw Bob lately and is going to be afraid to slap **now (9.45)**	I know <u>that</u> John said that Mary saw Bob lately and is going to be afraid to slap **us (4.21)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.40$		

(52)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John asked Mary to call lately and is going to be afraid to slap **now (10.08)**	*I know <u>who</u> John asked Mary to call lately and is going to be afraid to slap **us (3.59)** now
$-filler$	*I know <u>that</u> John asked Mary to call Bob lately and is going to be afraid to slap **now (9.83)**	I know <u>that</u> John asked Mary to call Bob lately and is going to be afraid to slap **us (4.72)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.38$		

(53)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John said that Mary saw recently and is going to be afraid to slap **now (10.12)**	*I know <u>who</u> John said that Mary saw recently and is going to be afraid to slap **us (3.78)** now
$-filler$	*I know <u>that</u> John said that Mary saw Bob recently and is going to be afraid to slap **now (9.98)**	I know <u>that</u> John said that Mary saw Bob recently and is going to be afraid to slap **us (5.01)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.37$		

B.10   ATB – Wikipedia LSTM

(54)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John asked Mary about recently and is going to kiss **soon (14.80)**	*I know <u>who</u> John asked Mary about recently and is going to kiss **us (8.39)** soon
$-filler$	*I know <u>that</u> John asked Mary about Bob recently and is going to kiss **soon (14.40)**	I know <u>that</u> John asked Mary about Bob recently and is going to kiss **us (11.63)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.64$		

(55)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John met recently and is going to beg Patricia to kiss **soon (15.45)**	*I know <u>who</u> John met recently and is going to beg Patricia to kiss **us (10.27)** soon
$-filler$	*I know <u>that</u> John met Bob recently and is going to beg Patricia to kiss **soon (15.46)**	I know <u>that</u> John met Bob recently and is going to beg Patricia to kiss **us (13.61)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.33$		

(56)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John asked Mary about recently and is going to be glad to kiss **soon (14.47)**	*I know <u>who</u> John asked Mary about recently and is going to be glad to kiss **us (9.19)** soon
$-filler$	*I know <u>that</u> John asked Mary about Bob recently and is going to be glad to kiss **soon (14.54)**	I know <u>that</u> John asked Mary about Bob recently and is going to be glad to kiss **us (12.59)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.33$		

(57)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John was happy to meet recently and is going to kiss **soon (13.61)**	*I know <u>who</u> John was happy to meet recently and is going to kiss **you (7.32)** soon
$-filler$	*I know <u>that</u> John was happy to meet Bob recently and is going to kiss **soon (13.36)**	I know <u>that</u> John was happy to meet Bob recently and is going to kiss **you (10.36)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.29$		

(58)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about recently and is going to beg Patricia to kiss **soon (15.49)**	*I know <u>who</u> John told Mary about recently and is going to beg Patricia to kiss **us (11.17)** soon
$-filler$	*I know <u>that</u> John told Mary about Bob recently and is going to beg Patricia to kiss **soon (15.26)**	I know <u>that</u> John told Mary about Bob recently and is going to beg Patricia to kiss **us (14.20)** soon
$\Delta_{-filler} - \Delta_{+filler} = -3.26$		

B.11   ATB – Wikipedia Transformer

(59)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John saw lately and is going to write to Patricia about **now (8.22)**	*I know <u>who</u> John saw lately and is going to write to Patricia about **us (7.18)** now
$-filler$	*I know <u>that</u> John saw Bob lately and is going to write to Patricia about **now (8.21)**	I know <u>that</u> John saw Bob lately and is going to write to Patricia about **us (8.38)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.21$		

(60)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John saw recently and is going to write to Patricia about **now (8.01)**	*I know <u>who</u> John saw recently and is going to write to Patricia about **us (7.28)** now
$-filler$	*I know <u>that</u> John saw Bob recently and is going to write to Patricia about **now (8.11)**	I know <u>that</u> John saw Bob recently and is going to write to Patricia about **us (8.53)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.15$		

(61)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John met lately and is going to write to Patricia about **now (8.26)**	*I know <u>who</u> John met lately and is going to write to Patricia about **us (7.40)** now
$-filler$	*I know <u>that</u> John met Bob lately and is going to write to Patricia about **now (8.35)**	I know <u>that</u> John met Bob lately and is going to write to Patricia about **us (8.58)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.09$		

(62)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John was happy to see lately and is going to write to Patricia about **now (8.46)**	*I know <u>who</u> John was happy to see lately and is going to write to Patricia about **us (7.52)** now
$-filler$	*I know <u>that</u> John was happy to see Bob lately and is going to write to Patricia about **now (8.45)**	I know <u>that</u> John was happy to see Bob lately and is going to write to Patricia about **us (8.54)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.04$		

(63)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John was happy to meet lately and is going to write to Patricia about **now (8.51)**	*I know <u>who</u> John was happy to meet lately and is going to write to Patricia about **us (7.68)** now
	$-filler$	*I know <u>that</u> John was happy to meet Bob lately and is going to write to Patricia about **now (8.57)**	I know <u>that</u> John was happy to meet Bob lately and is going to write to Patricia about **us (8.72)** now
	$\Delta_{-filler} - \Delta_{+filler} = -0.98$		

B.12 ATB – GPT-2

(64)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John said that Mary saw lately and is going to be glad to kiss **soon (14.91)**	*I know <u>who</u> John said that Mary saw lately and is going to be glad to kiss **you (4.06)** soon
	$-filler$	*I know <u>that</u> John said that Mary saw Bob lately and is going to be glad to kiss **soon (16.39)**	I know <u>that</u> John said that Mary saw Bob lately and is going to be glad to kiss **you (7.08)** soon
	$\Delta_{-filler} - \Delta_{+filler} = -1.53$		

(65)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John said that Mary met lately and is going to be glad to kiss **soon (14.46)**	*I know <u>who</u> John said that Mary met lately and is going to be glad to kiss **you (4.11)** soon
	$-filler$	*I know <u>that</u> John said that Mary met Bob lately and is going to be glad to kiss **soon (15.85)**	I know <u>that</u> John said that Mary met Bob lately and is going to be glad to kiss **you (6.93)** soon
	$\Delta_{-filler} - \Delta_{+filler} = -1.43$		

(66)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John said that Mary met lately and is going to be glad to kiss **now (12.15)**	*I know <u>who</u> John said that Mary met lately and is going to be glad to kiss **you (4.11)** now
	$-filler$	*I know <u>that</u> John said that Mary met Bob lately and is going to be glad to kiss **now (13.67)**	I know <u>that</u> John said that Mary met Bob lately and is going to be glad to kiss **you (6.93)** now
	$\Delta_{-filler} - \Delta_{+filler} = -1.29$		

(67)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John said that Mary met lately and is going to kiss **soon (13.26)**	*I know <u>who</u> John said that Mary met lately and is going to kiss **you (5.31)** soon
	$-filler$	*I know <u>that</u> John said that Mary met Bob lately and is going to kiss **soon (15.08)**	I know <u>that</u> John said that Mary met Bob lately and is going to kiss **you (8.30)** soon
	$\Delta_{-filler} - \Delta_{+filler} = -1.18$		

(68)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John was eager to see lately and is going to predict that Patricia will kiss **soon (11.50)**	*I know <u>who</u> John was eager to see lately and is going to predict that Patricia will kiss **Kim (10.84)** soon
	$-filler$	*I know <u>that</u> John was eager to see Bob lately and is going to predict that Patricia will kiss **soon (11.45)**	I know <u>that</u> John was eager to see Bob lately and is going to predict that Patricia will kiss **Kim (11.95)** soon
	$\Delta_{-filler} - \Delta_{+filler} = -1.16$		

B.13   ATB – GPT-j

(69)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John told Mary about lately and is going to hug **now (8.12)**	*I know <u>who</u> John told Mary about lately and is going to hug **us (6.59)** now
	$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to hug **now (8.88)**	I know <u>that</u> John told Mary about Bob lately and is going to hug **us (9.42)** now
	$\Delta_{-filler} - \Delta_{+filler} = -2.07$		

(70)		$+gap$	$-gap$
	$+filler$	I know <u>who</u> John told Mary about lately and is going to hug **today (8.69)**	*I know <u>who</u> John told Mary about lately and is going to hug **us (6.59)** today
	$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to hug **today (9.54)**	I know <u>that</u> John told Mary about Bob lately and is going to hug **us (9.42)** today
	$\Delta_{-filler} - \Delta_{+filler} = -1.98$		

(71)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John said that Mary met recently and is going to be afraid to hug **now (7.54)**	*I know <u>who</u> John said that Mary met recently and is going to be afraid to hug **us (4.59)** now
$-filler$	*I know <u>that</u> John said that Mary met Bob recently and is going to be afraid to hug **now (9.06)**	I know <u>that</u> John said that Mary met Bob recently and is going to be afraid to hug **us (8.07)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.96$		

(72)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John said that Mary met recently and is going to be afraid to slap **now (10.08)**	*I know <u>who</u> John said that Mary met recently and is going to slap **us (4.22)** now
$-filler$	*I know <u>that</u> John said that Mary met Bob recently and is going to be afraid to slap **now (10.79)**	I know <u>that</u> John said that Mary met Bob recently and is going to be afraid to slap **us (6.81)** now
$\Delta_{-filler} - \Delta_{+filler} = -1.88$		

(73)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about recently and is going to be afraid to hug **today (10.32)**	*I know <u>who</u> John told Mary about recently and is going to be afraid to hug **us (6.16)** today
$-filler$	*I know <u>that</u> John told Mary about Bob recently and is going to be afraid to hug **today (11.78)**	I know <u>that</u> John told Mary about Bob recently and is going to be afraid to hug **us (9.46)** today
$\Delta_{-filler} - \Delta_{+filler} = -1.84$		

B.14 ATB – GPT-3

(74)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about lately and is going to be afraid to hug **now (8.21)**	*I know <u>who</u> John told Mary about lately and is going to be afraid to hug **you (5.06)** now
$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to be afraid to hug **now (8.39)**	I know <u>that</u> John told Mary about Bob lately and is going to be afraid to hug **you (8.09)** now
$\Delta_{-filler} - \Delta_{+filler} = -2.85$		

(75)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about lately and is going to slap **now (9.43)**	*I know <u>who</u> John told Mary about lately and is going to slap **you (4.70)** now
$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to slap **now (10.25)**	I know <u>that</u> John told Mary about Bob lately and is going to slap **you (8.27)** now
$\Delta_{-filler} - \Delta_{+filler} = -2.75$		

(76)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about lately and is going to slap **soon (7.22)**	*I know <u>who</u> John told Mary about lately and is going to slap **you (4.70)** soon
$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to slap **soon (8.33)**	I know <u>that</u> John told Mary about Bob lately and is going to slap **you (8.27)** soon
$\Delta_{-filler} - \Delta_{+filler} = -2.47$		

(77)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about lately and is going to be glad to slap **soon (9.32)**	*I know <u>who</u> John told Mary about lately and is going to be glad to slap **you (3.35)** soon
$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to be glad to slap **soon (9.52)**	I know <u>that</u> John told Mary about Bob lately and is going to be glad to slap **you (5.82)** soon
$\Delta_{-filler} - \Delta_{+filler} = -2.27$		

(78)

	$+gap$	$-gap$
$+filler$	I know <u>who</u> John told Mary about lately and is going to be glad to slap **today (10.41)**	*I know <u>who</u> John told Mary about lately and is going to be glad to slap **you (3.35)** today
$-filler$	*I know <u>that</u> John told Mary about Bob lately and is going to be glad to slap **today (10.64)**	I know <u>that</u> John told Mary about Bob lately and is going to be glad to slap **you (5.82)** today
$\Delta_{-filler} - \Delta_{+filler} = -2.24$		

C Appendix: context-free grammars for retraining task (Section 5.3)

C.1 PG – retraining corpus

$\langle S \rangle \rightarrow \langle S\_FG \rangle$
$\langle S\_FG \rangle \rightarrow \langle PREAMBLE \rangle \ \langle F \rangle \ \langle G \rangle$
$\langle S\_XG \rangle \rightarrow \langle UNGRAMMATICAL \rangle \ \langle PREAMBLE \rangle \ \langle XF \rangle \ \langle G \rangle$
$\langle S\_FX \rangle \rightarrow \langle UNGRAMMATICAL \rangle \ \langle PREAMBLE \rangle \ \langle F \rangle \ \langle XG \rangle$
$\langle S\_XX \rangle \rightarrow \langle PREAMBLE \rangle \ \langle XF \rangle \ \langle XG \rangle$
$\langle GEN \rangle \rightarrow$ ''s'
$\langle OBJ \rangle \rightarrow$ 'you' | 'us' | 'Kim'
$\langle NAME1 \rangle \rightarrow$ 'Bob' | 'John'
$\langle NAME2 \rangle \rightarrow$ 'Mary' | 'Jennifer'
$\langle NAME3 \rangle \rightarrow$ 'James' | 'Michael'
$\langle NAME4 \rangle \rightarrow$ 'Patricia' | 'Linda'
$\langle PREAMBLE \rangle \rightarrow$ 'I know'
$\langle F \rangle \rightarrow$ 'who' $\langle NAME1 \rangle \ \langle GEN \rangle \ \langle SUBJ \rangle$
$\langle XF \rangle \rightarrow$ 'that' $\langle NAME1 \rangle \ \langle GEN \rangle \ \langle SUBJ \rangle \ \langle NAME2 \rangle$
$\langle G \rangle \rightarrow \langle G\_PAST \rangle \ | \ \langle G\_FUTURE \rangle$
$\langle XG \rangle \rightarrow \langle XG\_PAST \rangle \ | \ \langle XG\_FUTURE \rangle$
$\langle G\_PAST \rangle \rightarrow \langle VP\_PAST \rangle \ \langle ADV\_PAST \rangle$
$\langle G\_FUTURE \rangle \rightarrow \langle VP\_FUTURE \rangle \ \langle ADV\_FUTURE \rangle$
$\langle XG\_PAST \rangle \rightarrow \langle VP\_PAST \rangle \ \langle XG\_OBJ \rangle \ \langle ADV\_PAST \rangle$
$\langle XG\_FUTURE \rangle \rightarrow \langle VP\_FUTURE \rangle \ \langle XG\_OBJ \rangle \ \langle ADV\_FUTURE \rangle$
$\langle VP\_PAST \rangle \rightarrow$ 'upset' | 'distracted' | 'worried' | 'annoyed' | 'amused' | 'delighted'
$\langle VP\_FUTURE \rangle \rightarrow$ 'will' $\langle V\_FUTURE \rangle$
$\langle V\_FUTURE \rangle \rightarrow$ 'upset' | 'distract' | 'worry' | 'annoy' | 'amuse' | 'delight'
$\langle XG\_OBJ \rangle \rightarrow \langle NAME4 \rangle \ | \ \langle OBJ \rangle$
$\langle SUBJ \rangle \rightarrow$ 'attitude towards' | 'friendship with' | 'praising of' | 'fight with' | 'kissing with' | 'asking about'
$\langle ADV\_PAST \rangle \rightarrow$ 'yesterday' | 'recently' | 'often' | 'constantly' | 'today' | 'lately' | 'earlier'
$\langle ADV\_FUTURE \rangle \rightarrow$ 'today' | 'soon' | 'tomorrow' | 'now' | 'quickly'

C.2 ATB – retraining corpus

$\langle S \rangle \rightarrow \langle S\_FG \rangle$
$\langle S\_FG \rangle \rightarrow \langle PREAMBLE \rangle \ \langle F \rangle \ \langle G \rangle$
$\langle S\_XG \rangle \rightarrow \langle UNGRAMMATICAL \rangle \ \langle PREAMBLE \rangle \ \langle XF \rangle \ \langle G \rangle$
$\langle S\_FX \rangle \rightarrow \langle UNGRAMMATICAL \rangle \ \langle PREAMBLE \rangle \ \langle F \rangle \ \langle XG \rangle$
$\langle S\_XX \rangle \rightarrow \langle PREAMBLE \rangle \ \langle XF \rangle \ \langle XG \rangle$
$\langle GEN \rangle \rightarrow$ ''s'

$\langle OBJ \rangle \rightarrow$ 'you' | 'us' | 'Kim'

$\langle NAME1 \rangle \rightarrow$ 'John'

$\langle NAME2 \rangle \rightarrow$ 'Mary'

$\langle NAME3 \rangle \rightarrow$ 'Bob'

$\langle NAME4 \rangle \rightarrow$ 'Patricia'

$\langle PREAMBLE \rangle \rightarrow$ 'I know'

$\langle F \rangle \rightarrow$ 'who' $\langle SUBJ\_F \rangle$

$\langle XF \rangle \rightarrow$ 'that' $\langle SUBJ\_XF \rangle$

$\langle CONN \rangle \rightarrow$ 'and'

$\langle SUBJ\_F \rangle \rightarrow \langle ONE\_SUBJ\_F \rangle \mid \langle TWO\_SUBJ\_F \rangle$

$\langle SUBJ\_XF \rangle \rightarrow \langle ONE\_SUBJ\_XF \rangle \mid \langle TWO\_SUBJ\_XF \rangle$

$\langle ONE\_SUBJ\_F \rangle \rightarrow \langle NAME1 \rangle \; \langle V1 \rangle \; \langle ADV\_PAST\_1 \rangle \; \langle CONN \rangle$

$\langle TWO\_SUBJ\_F \rangle \rightarrow \langle NAME1 \rangle \; \langle V1 \rangle \; \langle ADV\_PAST\_1 \rangle \; \langle CONN \rangle \; \langle NAME3 \rangle$

$\langle ONE\_SUBJ\_XF \rangle \rightarrow \langle NAME1 \rangle \; \langle V1 \rangle \; \langle NAME2 \rangle \; \langle ADV\_PAST\_1 \rangle \; \langle CONN \rangle$

$\langle TWO\_SUBJ\_XF \rangle \rightarrow \langle NAME1 \rangle \; \langle V1 \rangle \; \langle NAME2 \rangle \; \langle ADV\_PAST\_1 \rangle \; \langle CONN \rangle \; \langle NAME3 \rangle$

$\langle G \rangle \rightarrow \langle G\_PAST \rangle \mid \langle G\_FUTURE \rangle$

$\langle XG \rangle \rightarrow \langle XG\_PAST \rangle \mid \langle XG\_FUTURE \rangle$

$\langle XG\_PAST \rangle \rightarrow \langle VP2\_PAST \rangle \; \langle XG\_OBJ \rangle \; \langle ADV\_PAST\_2 \rangle$

$\langle XG\_FUTURE \rangle \rightarrow \langle VP2\_FUTURE \rangle \; \langle XG\_OBJ \rangle \; \langle ADV\_FUTURE \rangle$

$\langle G\_PAST \rangle \rightarrow \langle VP2\_PAST \rangle \; \langle ADV\_PAST\_2 \rangle$

$\langle G\_FUTURE \rangle \rightarrow \langle VP2\_FUTURE \rangle \; \langle ADV\_FUTURE \rangle$

$\langle XG\_OBJ \rangle \rightarrow \langle NAME4 \rangle \mid \langle OBJ \rangle$

$\langle V1 \rangle \rightarrow$ 'saw' | 'hugged' | 'helped' | 'met' | 'pushed' | 'praised' | 'chased' | 'hired' | 'invited' | 'promoted' | 'w

$\langle VP2 \rangle \rightarrow \langle VP2\_PAST \rangle \mid \langle VP2\_FUTURE \rangle$

$\langle VP2\_PAST \rangle \rightarrow \langle V2\_PAST \rangle$

$\langle V2\_PAST \rangle \rightarrow$ 'kissed' | 'slapped' | 'insulted' | 'annoyed' | 'hurt' | 'mocked' | 'teased' | 'supported' | 'marrie

$\langle VP2\_FUTURE \rangle \rightarrow$ 'will' $\langle V2\_FUTURE \rangle$

$\langle V2\_FUTURE \rangle \rightarrow$ 'kiss' | 'slap' | 'insult' | 'annoy' | 'hurt' | 'mock' | 'tease' | 'support' | 'marry'

$\langle ADV\_PAST\_1 \rangle \rightarrow$ 'yesterday' | 'recently' | 'often' | 'constantly'

$\langle ADV\_PAST\_2 \rangle \rightarrow$ 'today' | 'lately' | 'earlier' | 'regularly' | 'repeatedly'

$\langle ADV\_FUTURE \rangle \rightarrow$ 'today' | 'soon' | 'tomorrow' | 'now' | 'quickly'