

Subword models struggle with word learning, but surprisal hides it

Bastian Bunzeck and Sina Zarriß

Computational Linguistics, Department of Linguistics

Bielefeld University, Germany

{bastian.bunzeck, sina.zarriess}@uni-bielefeld.de

Abstract

We study word learning in subword and character language models with the psycholinguistic lexical decision task. While subword LMs struggle to discern words and non-words with high accuracy, character LMs solve this task easily and consistently. Only when supplied with further contexts do subword LMs perform similarly to character models. Additionally, when looking at word-level and syntactic learning trajectories, we find that both processes are separable in character LMs. Word learning happens before syntactic learning, whereas both occur simultaneously in subword LMs. This raises questions about the adequacy of subword LMs for modeling language acquisition and positions character LMs as a viable alternative to study processes below the syntactic level.

1 Introduction

When humans acquire their first language(s), they *first* learn to recognize single words, mostly from short, fragmentary utterances (Cameron-Faulkner et al., 2003; Bunzeck and Diessel, 2024), *before* fully understanding the grammatical processes governing them (Tomasello, 1992; Behrens, 2021). This simple fact about language acquisition has received surprisingly little attention in the body of work that treats LMs as models of language learners (Warstadt and Bowman, 2022; Portelance and Jasbi, 2024). While word learning in children is comparatively well studied (Plunkett, 1997; Yu and Ballard, 2007; Waxman and Gelman, 2009; Bergelson and Swingle, 2012; Clark and Casillas, 2015; Frank et al., 2021), the implicit word learning processes in LMs are not. Current studies focus on syntax (Mueller et al., 2022; Choshen et al., 2022), or investigate word learning in close connection with syntax through surprisal (Chang and Bergen, 2022; Portelance et al., 2023; Shafabadi and Wisniewski, 2025; Ficarra et al., 2025). Architecture-wise, a key limitation to the precise study of word learning

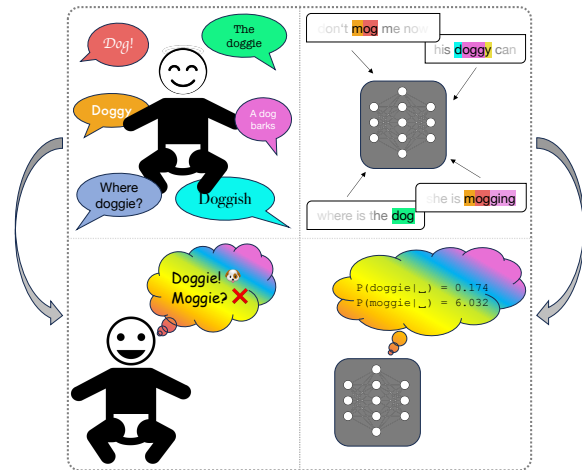


Figure 1: Illustration of word learning in human learners and transformer LLMs (top), and of our lexical decision test that probes discrimination of words from non-words (bottom). While human learners build up an mental lexicon from experience with language, artificial learners assign probabilities to strings based on their frequency.

is subword tokenization (e.g., BPE, Gage, 1994), which splits words into linguistically (Arnett and Bergen, 2025) and cognitively implausible units (Beinborn and Pinter, 2023).

To gauge word learning in a syntax-independent manner, we use the psycholinguistic lexical decision task (Meyer and Schvaneveldt, 1971; Le Goudais et al., 2017), i.e., deciding which word in a given word/non-word pair is real. We find that models with character-level tokenization learn this task quickly and reliably. In contrast, subword LMs of all sizes perform significantly worse in a syntax-independent setting and only achieve comparable accuracy when stimuli are measured through surprisal, or “unexpectedness,” in linguistic context. By comparing word and syntactic learning (measured via BLiMP, Warstadt et al., 2020), we further find that character models quickly acquire word knowledge and only later develop syntactic knowledge. In subword models, word and syntax learning happen concurrently. This demonstrates how ele-

mentary modeling decisions, such as tokenization methods, significantly impact learning trajectories in LMs, a fact that warrants more scrutiny when using LMs as models of language acquisition.

2 Related work

Word learning in humans is a multifaceted phenomenon that involves different kinds of (extra)linguistic knowledge (Waxman and Gelman, 2009): while phoneticians are concerned with word recognition and sequence segmentation (Jusczyk, 1999), developmental psychologists frequently equate word learning with correct reference to real world objects (Stager and Werker, 1997; Ackermann et al., 2020). Psycholinguists usually focus on the mental lexicon and learning which words belong to it (Goldinger, 1996), whereas usage-based scholars take into account children’s productions and their ability to be competent language users, even with few words (Tomasello, 1992).

Although aspects like word recognition and sequence segmentation have been studied in LMs (e.g. Goriely and Buttery, 2025a), the most common approach to word learning in LMs is measuring the predictability of words via surprisal (negative log-probability, Hale, 2001). Chang and Bergen (2022) train LMs on book texts and wiki data. They define a surprisal threshold below which words are said to be learned and find that frequent function words are learned earliest. Here, the alignment between models and real learners is questionable: children first utter nouns and verbs (Tomasello, 2000), but also rely on function words for challenges like speech segmentation (Dye et al., 2019). Portelance et al. (2023) show that in LSTMs trained on child-directed speech, surprisal correlates with word-level age of acquisition. Chang et al. (2024) observe that learning curves for surprisal values are stable for frequent tokens, while infrequent tokens are “forgotten” again over pre-training. Shafabadi and Wisniewski (2025) introduce anti-surprisal (in incorrect contexts) to track false usage, which also fluctuates over pre-training. These studies cast word learning as the ability to anticipate words’ expectedness in a given syntactic (and semantic) context. We note a certain conceptual leap to the original works on surprisal, where it is primarily viewed as an incremental measure of processing difficulty in syntactic comprehension (Levy, 2008; Demberg and Keller, 2009). A simple word like *dog* might be surprising and therefore hard to parse

in some contexts, but very expected in others, independently of being already learned on the pure word level. A further methodological drawback of surprisal as a measure of word learning is that it corresponds almost directly to the next-token prediction objective LMs are trained on. This contrasts with typical probing paradigms used in the domain of syntax, which implement the idea to “challenge” models in minimal pair set-ups that are not observed directly as string sequences in training, thereby testing abstracted, implicit linguistic knowledge rather than observed patterns in the data. In a similar vein, we want to probe the word knowledge of an LM at a fundamental level and beyond surface-level word sequences that LMs are known to excel in predicting. We want to know if the artificial learner knows that the word *doggie* exists in the English language, but *moggie* does not.

Lexical decision is widely used in human studies but remains an underexplored LM benchmark. Le Godais et al. (2017) show that character-based LSTMs achieve about 95% accuracy on such tasks. Lavechin et al. (2023) find that speech-based LMs need significantly more input to still perform poorly (56.8%) than phoneme-level LSTMs (75.4%) on a phonetic dataset. For the same data, Goriely et al. (2024) find that GPT-2-based subword BabyLMs achieve 70% accuracy and a comparable character models reach nearly 90%. Finally, for another lexical decision dataset, Bunzeck et al. (2025) report near-perfect accuracy for character-based grapheme Llama models, while phoneme models perform at 60–70%.

3 Experiments

Models We train triplets of increasingly larger Llama models (Touvron et al., 2023) with character/subword tokenization on the BabyLM 10M corpus (Choshen et al., 2024). Training details are found in Appendix A. As ablations, we test subword Pythias (Biderman et al., 2023) and character/subword GPT-2 models (Goriely et al., 2024).

Test data We follow the idea of forced-choice lexical decision (Baddeley et al., 1993), where participants must decide which is real: an existing word or a synthesized non-word. We use wuggy (Keuleers and Brysbaert, 2010) to generate minimal pairs of words/non-words that differ in one or two syllables, akin to syntactic minimal pair tests such as BLIMP. We derive 1,000 non-words (e.g. *mondning*) each from 1,000 high-frequency/low-

Tokenization	Model	Parameters	Data size	Lexical decision		Surprisal		Anti-surprisal	
				highFrq	lowFrq	highFrq	lowFrq	highFrq	lowFrq
Subword (BPE)	Pythia	14M	825GB	66.6	62.5	90.5	85.5	71.4	77.7
		70M		72.5	68.8	94.5	94.0	77.0	83.6
		160M		77.8	73.0	96.4	95.8	78.0	85.7
		410M		81.9	78.1	97.7	97.9	77.1	84.1
		1B		87.5	83.2	97.7	97.9	76.6	83.8
		1.4B		87.8	81.6	97.9	97.9	76.5	84.7
	GPT-2	97.5M	100M words	35.6	79.1	99.0	99.2	84.7	86.9
	Llama	2.51M	10M words	70.9	58.4	86.7	70.9	78.6	67.7
		7.77M		79.5	63.2	91.3	78.1	81.1	72.9
		30.03M		83.6	68.6	92.7	81.1	83.7	76.1
Character	GPT-2	85.3M	100M words	98.7	97.3	99.8	99.4	98.0	96.3
	Llama	0.49M	10M words	97.6	83.0	98.2	84.3	98.0	83.1
		3.73M		98.9	90.2	99.4	90.3	98.5	88.8
		21.94M		99.0	93.3	99.8	94.7	99.0	92.5

Table 1: Accuracy scores (in %) for (i) lexical decision, (ii) surprisal and (iii) anti-surprisal experiments

frequency words (e.g. *sending*), which preserve syllable-bigram frequencies and match their origin words in length (cf. Appendix B).

Lexical decision For a word/non-word pair $(w, *w)$, we measure $-\log(P(w|_))$ and $-\log(P(*w|_))$, i.e. how “surprised” a LM is by the word in the context of a prepended whitespace (and BOS token). If $-\log(P(w|_)) < -\log(P(*w|_))$, the LM’s lexical decision is correct. As autoregressive LMs are sequence prediction models, we need a preceding context for which we can calculate surprisal. A single whitespace is the most neutral starting token available (and for subword models also signals that the first subword is word-initial). For all experiments, we calculate the average surprisal over all tokens of a word (which, in some cases, is characterized by a mismatch in token numbers between words and non-words, cf. Appendix B) with minicons (Misra, 2022).

Surprisal To measure LMs’ knowledge of words presented in regular syntactic contexts, we calculate the surprisal of words and non-words $(w, *w)$ as $-\log(P(w_i|w_{n<i}))$, i.e. the degree to which the LM is “surprised” by the word in the context of plausible preceding tokens, including a BOS token. We create stimuli by sampling sentences that contain our target words from OpenSubtitles (Lison and Tiedemann, 2016) and substituting them with matching non-words for the false stimuli. If $-\log(P(w_i|w_{n<i})) < -\log(P(*w_i|w_{n<i}))$, the LM’s decision is correct.

Anti-surprisal Inspired by Shafiabadi and Wisniewski (2025), we include anti-surprisal, a mea-

sure of word surprisal on negative instances. We create negative samples by selecting sentences that our original words do not occur in, and then randomly¹ placing words/non-words into these sentences at the same index, where $index \geq 3$. By doing so, we compromise between lexical decision and surprisal measurement. There are two reasons to include this measure: i) surprisal in negative samples provides the model with word material as context, but without semantic or syntactic signals that could prime the model towards recognizing it; this allows us to assess whether the mere presence of other words in context makes it easier for the model to distinguish words from non-words, compared to our lexical decision set-up where only whitespace is given. In addition, ii) we want to see if the presence of an ill-fitting context actively deteriorates performance in the sense of a model suddenly preferring non-words over existing words. Again, if $-\log(P(w_i|w_{n<i})) < -\log(P(*w_i|w_{n<i}))$, the LM’s decision is correct.

Learning trajectories To assess when word learning happens in relation to syntax learning, we further evaluate intermediate checkpoints of our models on our word learning tests and BLiMP as a syntactic benchmark. In line with previous studies (Chang and Bergen, 2022; Viering and Loog, 2023), we space our checkpoints logarithmically – 10 for the first 10% of training, 9 more for the remaining 90%. For the Pythia models, we extract similarly spaced checkpoints (the GPT-2 models do not provide checkpoints, so we exclude them).

¹In line with Shafiabadi and Wisniewski (2025), we do not further characterize the resulting violations on a syntactic or semantic level.

4 Results

Lexical decision The lexical decision results (Table 1) show a strong contrast between character and subword models. Character models achieve near-perfect accuracy (97–99%) on high-frequency words, regardless of model size. The performance on low-frequency words steadily increases with model size and reaches a near-perfect level for our largest character Llama and the character GPT-2. On the other hand, all BPE models get surprisingly low scores on high-frequency words: The smallest Pythia model discriminates between word and non-words with an accuracy of only 67%, and the BPE GPT-2 performs below the chance baseline. Even the largest and best BPE model reaches only 87.8% on high-frequency words – almost 10% less than the smallest character model.

Scaling laws generally hold, with larger models outperforming smaller ones. Interestingly, for the BPE models, there is a consistent gap between high and low-frequency words that cannot be closed by larger models. Smaller character models also show a performance gap between high and low-frequency words, but it narrows considerably with larger models. These results point to substantial differences in how subword and character models learn words. Such a surprising lack of ability in distinguishing words from non-words (without context) is a blatant, hitherto overlooked gap in subword models.

Surprisal and anti-surprisal Results for the second experiment (Table 1) differ from those for lexical decision. In the surprisal setting, the difference between BPE and character models is less pronounced. On high-frequency data, nearly all models (except the smallest BPE Llama) achieve over 90% accuracy. Still, larger character models yield the best results. In the low-frequency data condition, the pattern is similar, though scores are generally lower. Very large BPE models outperform our Llamas there, but the character GPT-2 remains superior. This may be attributed to the limited lexical exposure of our Llamas, trained on only 10M tokens. In the anti-surprisal setting, character models again drastically outperform BPE models and achieve nearly perfect scores on high-frequency data, while BPE models only reach 70–80% accuracy. This setting is the only one where Pythia models get better scores on low-frequency data, with a gap of 6–8%, which increases with model size (not the case for BPE Llamas or character models). This contrast between surprisal and antisurprisal might

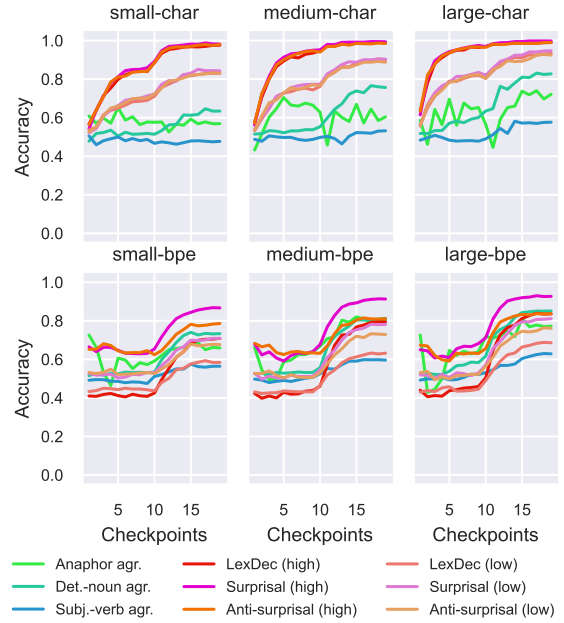


Figure 2: Selected lexical and syntactic learning curves

be an indicator of the entanglement of word learning and syntactic learning in subword models. It is plausible that for high-frequency words, the BPE models have strong expectations about which word should come next in a certain context, and because this expectation is not matched by the real (but ill-fitting) word, a made-up non-word is preferred. We argue that this should still not be the case for an ideal language model – if a model is indeed well-tuned, it should assign a higher probability to an ill-fitting but existing word which is still in distribution, than to a completely ill-fitting string that is out-of-distribution. In any case, BPE models catch up to character models if (and *only* if) provided with additional syntactic/semantic context information. While random context somewhat aids BPE models, a substantial gap remains between the largest BPE models and the character models, where performance remains excellent, even in the presence of implausible contexts.

Learning trajectories Figure 2 displays learning curves for syntactic agreement phenomena, lexical decision, and both surprisal conditions across the 19 saved checkpoints (complete curves reproduced in Appendix C). The first 10 checkpoints correspond to the first 10% of pretraining, the remaining 9 checkpoints represent 10% of training steps over the remaining 90% of pretraining. For character models, the high-frequency, low-frequency, and syntactic curves are clearly separated. On high-frequency data, word learning is rapid and follows power-law curves; the low-frequency scores im-

prove a little later and at a lower rate, but with the same trajectory. Syntactic phenomena improve later, mostly in s-shaped curves (e.g., det.-noun agreement), or are not learned at all in small models (subj.-verb agreement). In contrast, the syntactic and lexical curves for BPE models form sheaves of s-shaped trajectories. There is *no* principal difference in learning dynamics between the syntactic and the word level, improvements² occur simultaneously. This further confirms the results of our previous experiments: in BPE models, word learning is dependent on syntax learning, and words cannot be recognized reliably outside of plausible contexts. Additionally, these different levels of learning cannot be disentangled in BPE models, whereas in character models, syntactic learning *follows* word learning.³

5 Discussion

In contrast to previous work, this study set out to explore when (and if) LMs learn what valid words *are*, and not how LMs learn when words are validly *used*. How should these results now be interpreted in the light of language acquisition and the use of BabyLMs to model corresponding processes? In reality, words are not learned in isolation, but from usage. Yet, words have been widely shown to be represented as solid “standalone” units in the mental lexicon, and to be units that human learners acquire early on (cf. Waxman and Gelman, 2009, also Montag et al., 2018). Usage-based approaches have tested aspects of word learning independently from syntax (mostly in object naming tasks, cf. Tomasello and Todd, 1983) and relate it to pragmatic aspects of communication. For example, children only react to words they know and ignore similar-sounding words. They even struggle with learning words that are phonetically extremely similar (like our stimuli), and only later gain this capacity (Stager and Werker, 1997). Similarly, in production, the earliest words come in isolation, slowly emerge into pivot schemas and holophrases (Tomasello, 1992, 2003), and only then finally turn into complex sentences. Of course, syntax also aids in discovering aspects of “wordiness”, like SV(X)-sentences offering cues for agent-patient relationships. It would be very interesting to further disentangle these levels of word knowledge

²We calculate Spearman-rank correlation between ordered accuracies for word-level tasks and BLiMP in Appendix D.

³Interestingly, the learning process is not finished when accuracy curves stabilize, cf. Appendix F.

in a follow-up study, but our current study is not concerned with the “you shall know a word by the company it keeps”-level of word knowledge, but rather with the “what do valid words of the language look like”-level, which can be assessed via lexical decision. As such, we believe that the separated learning curves of our character-level models represent more human-like learning than the highly correlated curves found in the subword models, but in reality, an overlap between them is definitely expected.

Reasons for the tremendous performance differences between subword and character models remain open to further inquiry. One plausible explanation is that character models have much more context available to calculate meaningful sequence probabilities, as words are split into many more tokens. While this is true, it is also exactly the point that we are stressing here: it is hard to imagine that arbitrary subword units lead to human-like, plausible word-level representations (like in exemplar models of lexical storage, cf. Bybee, 2010), whereas character models might offer a more justified level of granularity (and, e.g., better fit reading times, cf. Oh et al., 2021). Our findings also align with results on LMs’ sensitivity to character-level perturbations (Moradi and Samwald, 2021; Zhu et al., 2024) and their inability to solve character-level tasks, like counting occurrences of the letter *r* in the word *strawberry* (Zhang and He, 2024; Shin and Kaneko, 2024; Cosma et al., 2025).

6 Conclusion

We have shown that the lexical decision approach to the study of word learning in LMs complements surprisal-based approaches and reveals difficulties that surprisal hides: subword LMs struggle with lexical decision, whereas character models master this task with ease. Additionally, in subword LMs, lexical and syntactic learning are inseparable, whereas word learning in character models precedes syntactic learning; the processes are related, yet separable. It is plausible that the *a priori* token splitting in subword models preempts a word discovery process in them, whereas character models first have to pass through this developmental stage, possibly in a somewhat more human-like manner. In any case, as we have shown, decisions about the representational levels of LMs tremendously influence their learning pathways on the different levels of linguistic analysis.

Limitations

The generalizability of our findings is constrained by a few factors. The present study has only focused on the English language, but it is plausible that other languages with different writing systems or graphematic and phonotactic rule systems exhibit different patterns under different tokenization schemes. Here, phonetic transcriptions might provide a viable alternative, but real narrow transcriptions that accurately capture the whole breadth of human input are scarce and extremely costly and laborious to manually produce (although novel datasets like [Goriely and Buttery, 2025b](#) provide an alternative through automatic transcription). Besides, for the character LMs, we focus only on small models, as very large models with such tokenization, especially ones providing intermediate checkpoints, are nonexistent at this moment. It would still be interesting to see how they compare to subword models in a setting where parameter size and training data are greatly increased.

As already mentioned in Section 2, from a developmental perspective, word learning in humans also includes other processes than statistical pattern recognition from the input: semantic aspects and real-world reference are equally important, as are multimodal input and communicative intent. The ongoing form-vs.-function debate on LMs ([Machowald et al., 2024](#)) has begun to consider these aspects, and further studies should aim at incorporating them; for example, the object naming paradigm used in many developmental psychology studies would lend itself naturally to the study of multimodal models.

Finally, we also want to mention that there are attempts to add more linguistic theory to tokenizers. Looking into different tokenizers that, for example, try to be more morphology-aware ([Hofmann et al., 2022](#); [Bauwens and Delobelle, 2024](#); [Yehezkel and Pinter, 2023](#); [Uzan et al., 2024](#)) or implement other optimization tricks ([Schmidt et al., 2024](#)) could yield even more fine-grained points of comparison, but for the present study and its limited scope we focused on the most popular tokenization scheme (BPE) and the most linguistically minimalist alternative (characters).

Ethical considerations

Due to the nature of this work, no concrete ethical aspects or repercussions need to be discussed. However, we would like to stress that, of course,

BabyLMs not supposed to simulate real babies, but only abstractions of a very specific part of their learning capacity (frequency-driven, domain-general learning mechanisms such as entrenchment or resonance), and therefore all claims about their implications for language development in the real world should be interpreted in this light.

Acknowledgments

We would like to thank the three anonymous ARR reviewers for their insightful comments and the highly valuable discussions throughout the rebuttal period.

This research has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – CRC-1646, project number 512393437, project A02.

References

- Lena Ackermann, Robert Hepach, and Nivedita Mani. 2020. [Children learn words easier when they are interested in the category to which the word belongs](#). *Developmental Science*, 23(3):e12915.
- Philip A. Allen, Albert F. Smith, Mei-Ching Lien, Jeremy Grabbe, and Martin D. Murphy. 2005. [Evidence for an Activation Locus of the Word-Frequency Effect in Lexical Decision](#). *Journal of Experimental Psychology: Human Perception and Performance*, 31(4):713–721.
- Catherine Arnett and Benjamin Bergen. 2025. [Why do language models perform worse for morphologically complex languages?](#) In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6607–6623, Abu Dhabi, UAE. Association for Computational Linguistics.
- Harald Baayen, Piepenbrock R, and Gulikers L. 1995. [CELEX2](#).
- Alan Baddeley, Hazel Emslie, and Ian Nimmo-Smith. 1993. [The Spot-the-Word test: A robust estimate of verbal intelligence based on lexical decision](#). *British Journal of Clinical Psychology*, 32(1):55–65.
- Thomas Bauwens and Pieter Delobelle. 2024. [BPE-knockout: Pruning Pre-existing BPE Tokenisers with Backwards-compatible Morphological Semi-supervision](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5810–5832, Mexico City, Mexico. Association for Computational Linguistics.
- Heike Behrens. 2021. [Constructivist Approaches to First Language Acquisition](#). *Journal of Child Language*, 48(5):959–983.

- Lisa Beinborn and Yuval Pinter. 2023. [Analyzing Cognitive Plausibility of Subword Tokenization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4478–4486, Singapore. Association for Computational Linguistics.
- Elika Bergelson and Daniel Swingley. 2012. [At 6–9 months, human infants know the meanings of many common nouns](#). *Proceedings of the National Academy of Sciences*, 109(9):3253–3258.
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. [Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling](#). *Preprint*, arXiv:2304.01373.
- Bastian Bunzeck and Holger Diessel. 2024. [The richness of the stimulus: Constructional variation and development in child-directed speech](#). *First Language*, 45(2):152–176.
- Bastian Bunzeck, Daniel Duran, Leonie Schade, and Sina Zarriß. 2025. [Small language models also work with small vocabularies: Probing the linguistic abilities of grapheme- and phoneme-based baby llamas](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6039–6048, Abu Dhabi, UAE. Association for Computational Linguistics.
- Joan Bybee. 2010. *Language, Usage and Cognition*, 1 edition. Cambridge University Press.
- Thea Cameron-Faulkner, Elena Lieven, and Michael Tomasello. 2003. [A construction based analysis of child directed speech](#). *Cognitive Science*, 27(6):843–873.
- Tyler A. Chang and Benjamin K. Bergen. 2022. [Word Acquisition in Neural Language Models](#). *Transactions of the Association for Computational Linguistics*, 10:1–16.
- Tyler A. Chang, Zhuowen Tu, and Benjamin K. Bergen. 2024. [Characterizing Learning Curves During Language Model Pre-Training: Learning, Forgetting, and Stability](#). *Transactions of the Association for Computational Linguistics*, 12:1346–1362.
- Leshem Choshen, Ryan Cotterell, Michael Y. Hu, Tal Linzen, Aaron Mueller, Candace Ross, Alex Warstadt, Ethan Wilcox, Adina Williams, and Chengxu Zhuang. 2024. [\[Call for Papers\] The 2nd BabyLM Challenge: Sample-efficient pretraining on a developmentally plausible corpus](#). *Preprint*, arXiv:2404.06214.
- Leshem Choshen, Guy Hach Cohen, Daphna Weinshall, and Omri Abend. 2022. [The Grammar-Learning Trajectories of Neural Language Models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8281–8297, Dublin, Ireland. Association for Computational Linguistics.
- Eve V Clark and Marisa Casillas. 2015. First language acquisition. In *The Routledge Handbook of Linguistics*, pages 311–328. Routledge.
- Adrian Cosma, Stefan Ruseti, Emilian Radoi, and Mihai Dascalu. 2025. [The Strawberry Problem: Emergence of Character-level Understanding in Tokenized Language Models](#). *arXiv preprint*.
- Vera Demberg and Frank Keller. 2009. [A Computational Model of Prediction in Human Parsing: Unifying Locality and Surprisal Effects](#). In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 31, pages 1888–1893.
- Cristina Dye, Yarden Kedar, and Barbara Lust. 2019. [From lexical to functional categories: New foundations for the study of language development](#). *First Language*, 39(1):9–32.
- Filippo Ficara, Ryan Cotterell, and Alex Warstadt. 2025. [A Distributional Perspective on Word Learning in Neural Language Models](#). *arXiv preprint*.
- Michael C. Frank, Mika Braginsky, Daniel Yurovsky, and Virginia A. Marchman. 2021. *Variability and Consistency in Early Language Learning: The Wordbank Project*. The MIT Press.
- Philip Gage. 1994. [A new algorithm for data compression](#). *The C Users Journal*, 12(2):23–38.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The Pile: An 800GB Dataset of Diverse Text for Language Modeling](#). *Preprint*, arXiv:2101.00027.
- Stephen D. Goldinger. 1996. [Auditory Lexical Decision](#). *Language and Cognitive Processes*, 11(6):559–568.
- Zébulon Goriely and Paula Buttery. 2025a. [BabyLM’s First Words: Word Segmentation as a Phonological Probing Task](#). *arXiv preprint*.
- Zébulon Goriely and Paula Buttery. 2025b. [IPA-CHILDES & G2P+: Feature-Rich Resources for Cross-Lingual Phonology and Phonemic Language Modeling](#). *arXiv preprint*.
- Zébulon Goriely, Richard Diehl Martinez, Andrew Caines, Paula Buttery, and Lisa Beinborn. 2024. [From babble to words: Pre-training language models on continuous streams of phonemes](#). In *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pages 37–53, Miami, FL, USA. Association for Computational Linguistics.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.

- Valentin Hofmann, Hinrich Schuetze, and Janet Pierrehumbert. 2022. [An Embarrassingly Simple Method to Mitigate Undesirable Properties of Pretrained Language Model Tokenizers](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 385–393, Dublin, Ireland. Association for Computational Linguistics.
- Peter W. Jusczyk. 1999. [How infants begin to extract words from speech](#). *Trends in Cognitive Sciences*, 3(9):323–328.
- Emmanuel Keuleers and Marc Brysbaert. 2010. [Wuggy: A multilingual pseudoword generator](#). *Behavior Research Methods*, 42(3):627–633.
- Marvin Lavechin, Yaya Sy, Hadrien Titeux, María Andrea Cruz Blandón, Okko Räsänen, Hervé Bredin, Emmanuel Dupoux, and Alejandrina Cristia. 2023. [BabySLM: Language-acquisition-friendly benchmark of self-supervised spoken language models](#). In *INTERSPEECH 2023*, pages 4588–4592. ISCA.
- Gaël Le Godais, Tal Linzen, and Emmanuel Dupoux. 2017. [Comparing character-level neural language models using a lexical decision task](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 125–130, Valencia, Spain. Association for Computational Linguistics.
- Roger Levy. 2008. [Expectation-based syntactic comprehension](#). *Cognition*, 106(3):1126–1177.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Kyle Mahowald, Anna A. Ivanova, Idan A. Blank, Nancy Kanwisher, Joshua B. Tenenbaum, and Evelina Fedorenko. 2024. [Dissociating language and thought in large language models](#). *Trends in Cognitive Sciences*, pages 517–540.
- James L. McClelland and David E. Rumelhart. 1981. [An interactive activation model of context effects in letter perception: I. An account of basic findings](#). *Psychological Review*, 88(5):375–407.
- David E. Meyer and Roger W. Schvaneveldt. 1971. [Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations](#). *Journal of Experimental Psychology*, 90(2):227–234.
- Kanishka Misra. 2022. [Minicons: Enabling Flexible Behavioral and Representational Analyses of Transformer Language Models](#). *arXiv preprint*.
- Jessica L. Montag, Michael N. Jones, and Linda B. Smith. 2018. [Quantity and Diversity: Simulating Early Word Learning Environments](#). *Cognitive Science*, 42(S2):375–412.
- Milad Moradi and Matthias Samwald. 2021. [Evaluating the Robustness of Neural Language Models to Input Perturbations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1558–1570, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Aaron Mueller, Robert Frank, Tal Linzen, Luheng Wang, and Sebastian Schuster. 2022. [Coloring the Blank Slate: Pre-training Imparts a Hierarchical Inductive Bias to Sequence-to-sequence Models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1352–1368, Dublin, Ireland. Association for Computational Linguistics.
- Byung-Doh Oh, Christian Clark, and William Schuler. 2021. [Surprisal Estimators for Human Reading Times Need Character Models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3746–3757, Online. Association for Computational Linguistics.
- Kim Plunkett. 1997. [Theories of early language acquisition](#). *Trends in Cognitive Sciences*, 1(4):146–153.
- Eva Portelance, Yuguang Duan, Michael C. Frank, and Gary Lupyan. 2023. [Predicting Age of Acquisition for Children’s Early Vocabulary in Five Languages Using Language Model Surprisal](#). *Cognitive Science*, 47(9):e13334.
- Eva Portelance and Masoud Jasbi. 2024. [The Roles of Neural Networks in Language Acquisition](#). *Language and Linguistics Compass*, 18(6):e70001.
- Craig W Schmidt, Varshini Reddy, Haoran Zhang, Alec Alameddine, Omri Uzan, Yuval Pinter, and Chris Tanner. 2024. [Tokenization Is More Than Compression](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 678–702, Miami, Florida, USA. Association for Computational Linguistics.
- Nazanin Shafiabadi and Guillaume Wisniewski. 2025. [Beyond surprisal: A dual metric framework for lexical skill acquisition in LLMs](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6636–6641, Abu Dhabi, UAE. Association for Computational Linguistics.
- Andrew Shin and Kunitake Kaneko. 2024. [Large language models lack understanding of character composition of words](#). In *ICML 2024 Workshop on Lms and Cognition*.
- Christine L. Stager and Janet F. Werker. 1997. [Infants listen for more phonetic detail in speech perception than in word-learning tasks](#). *Nature*, 388(6640):381–382.
- Michael Tomasello. 1992. *First Verbs: A Case Study of Early Grammatical Development*, 1 edition. Cambridge University Press.

- Michael Tomasello. 2000. The item-based nature of children’s early syntactic development. *Trends in Cognitive Sciences*, 4(4).
- Michael Tomasello. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press.
- Michael Tomasello and Jody Todd. 1983. [Joint attention and lexical acquisition style](#). *First Language*, 4(12):197–211.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [LLaMA: Open and Efficient Foundation Language Models](#). *Preprint*, arXiv:2302.13971.
- Omri Uzan, Craig W. Schmidt, Chris Tanner, and Yuval Pinter. 2024. [Greed is All You Need: An Evaluation of Tokenizer Inference Methods](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 813–822, Bangkok, Thailand. Association for Computational Linguistics.
- Tom Viering and Marco Loog. 2023. [The Shape of Learning Curves: A Review](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7799–7819.
- Alex Warstadt and Samuel R. Bowman. 2022. *What Artificial Neural Networks Can Tell Us about Human Language Acquisition*, 1 edition, pages 17–60. CRC Press, Boca Raton.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Gotlieb Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, and Ryan Cotterell. 2023. [Findings of the BabyLM Challenge: Sample-Efficient Pretraining on Developmentally Plausible Corpora](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 1–6, Singapore. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The Benchmark of Linguistic Minimal Pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Sandra R. Waxman and Susan A. Gelman. 2009. [Early word-learning entails reference, not merely associations](#). *Trends in Cognitive Sciences*, 13(6):258–263.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick Von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Shaked Yehezkel and Yuval Pinter. 2023. [Incorporating context into subword vocabularies](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 623–635, Dubrovnik, Croatia. Association for Computational Linguistics.
- Chen Yu and Dana H. Ballard. 2007. [A unified model of early word learning: Integrating statistical and social cues](#). *Neurocomputing*, 70(13-15):2149–2165.
- Yidan Zhang and Zhenan He. 2024. [Large Language Models Can Not Perform Well in Understanding and Manipulating Natural Language at Both Character and Word Levels?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11826–11842, Miami, Florida, USA. Association for Computational Linguistics.
- Kaijie Zhu, Qinlin Zhao, Hao Chen, Jindong Wang, and Xing Xie. 2024. [PromptBench: A unified library for evaluation of large language models](#). *Journal of Machine Learning Research*, 25(254):1–22.

A Model hyperparameters and training details

We use the transformers library (Wolf et al., 2020) to train our models. The corresponding hyperparameters are listed in Table 2. We opted for very small LMs and little training data because the BabyLM paradigm has consolidated itself as a well-established method of developmentally plausible language modeling (Warstadt et al., 2023; Choshen et al., 2024). Also, we noticed the word learning to occur quite rapidly in our models, so we argue that small LMs offer more fine-grained opportunities for investigating these processes – in larger LMs, even a singular training step can already influence performance on such a brittle task like lexical decision tremendously.

The subword models feature a considerably higher number of parameters, as the embedding layer of transformer language models accounts for a quite large share of overall model parameters. In light of our vastly different vocabulary sizes (102 for character models, 8,002 for subword models), these differences are not surprising. The subword tokenizer is a regular BPE tokenizer self-trained with the tokenizers library and contains 8,000 subword tokens, a beginning-of-sequence token and a combined end-of-sequence/padding token. The character tokenizer contains these two special tokens and all printable ASCII characters, which are sufficient to represent all graphemes of the English language.

Our models were trained on a Apple M2 Pro processor with the MPS backend. For the small models, training lasted approx. 20min, for the medium-sized models it took approx. 1h and for the largest models approx. 10h (no principal differences between character and subword models). We share our models with their checkpoints on the HuggingFace hub.⁴

Loss curves for all models can be found in Figure 3. For the test loss, we calculated the perplexity over a held-out portion of our training corpus that is comparable in composition to the training data. We find no principal differences in loss development, although the character models converge faster. Larger models also tend to converge faster and generally reach smaller absolute loss values. As the similar train and test loss curves indicate, all Llama models succeed in optimizing for their next-

token prediction objective. However, it remains open to further inquiry how much these scores are constrained by the comparatively small capacity of our BabyLMs, and whether larger models would enhance performance further.

B Data creation and tokenization analysis

Data creation Word frequency influences lexical decision performance greatly (McClelland and Rumelhart, 1981; Allen et al., 2005). To incorporate this effect into our study, we create two distinct data sets from words included in wuggy: (i) high-frequency stimuli with a frequency score over 7.0 (occurrences per 1M words in BNC, COCA and other English corpora, as reported in CELEX, Baayen et al., 1995) and (ii) low-frequency stimuli with a frequency score over 0.0 but below 0.7 (so at least one order of magnitude lower). We opted to rely on the CELEX frequency scores because we compare models trained on different corpora – our models are trained on the 10M BabyLM corpus, the models by Goriely et al. (2024) are trained on the 100M BabyLM corpus and the Pythia models are trained on The Pile (Gao et al., 2020). As such, frequency scores from these corpora would taint analyses of other models and hinder comparability. For the contextualised stimuli, we sample sentences from the OpenSubtitles (Lison and Tiedemann, 2016) portion of the BabyLM 2024 corpus (Choshen et al., 2024). Both Wuggy and BabyLM data are licensed under the MIT license⁵, therefore we release our own stimuli artifacts on HuggingFace⁶ under the same license. The data contain no information that names or uniquely identifies individual people or offensive content, and are commonly used in computational linguistics.

Analysis of tokenization To further assess the influence that these frequency scores have on the resulting tokenization for our own models, we offer a brief analysis: Figure 4 shows a pairplot between three numerical variables – (i) the number of tokens that our original words are split into, (ii) the number of tokens that the derived non-word are split into and (iii) the corresponding frequency score from CELEX. While the three plots on the diagonal axis show a layered kernel density estimate (KDE) for

⁴<https://huggingface.co/collections/bbunzeck/word-learning-in-small-lms-67bdc21868856162b3be08f>

⁵As per license information found at <https://github.com/WuggyCode/wuggy> and <https://github.com/babylm/evaluation-pipeline-2024>.

⁶<https://huggingface.co/datasets/bbunzeck/lexical-decision>

	small-char	medium-char	large-char	small-bpe	medium-bpe	large-bpe
Embedding size	128	256	512	128	256	512
Hidden size	128	256	512	128	256	512
Layers	4	8	12	4	8	12
Attention heads	4	8	12	4	8	12
Context size	128	128	128	128	128	128
Vocab. size	102	102	102	8,002	8,002	8,002
Parameters	486,016	3,726,592	21,940,736	2,508,416	7,771,392	30,030,336

Table 2: Model hyperparameters for our self-trained Llama models

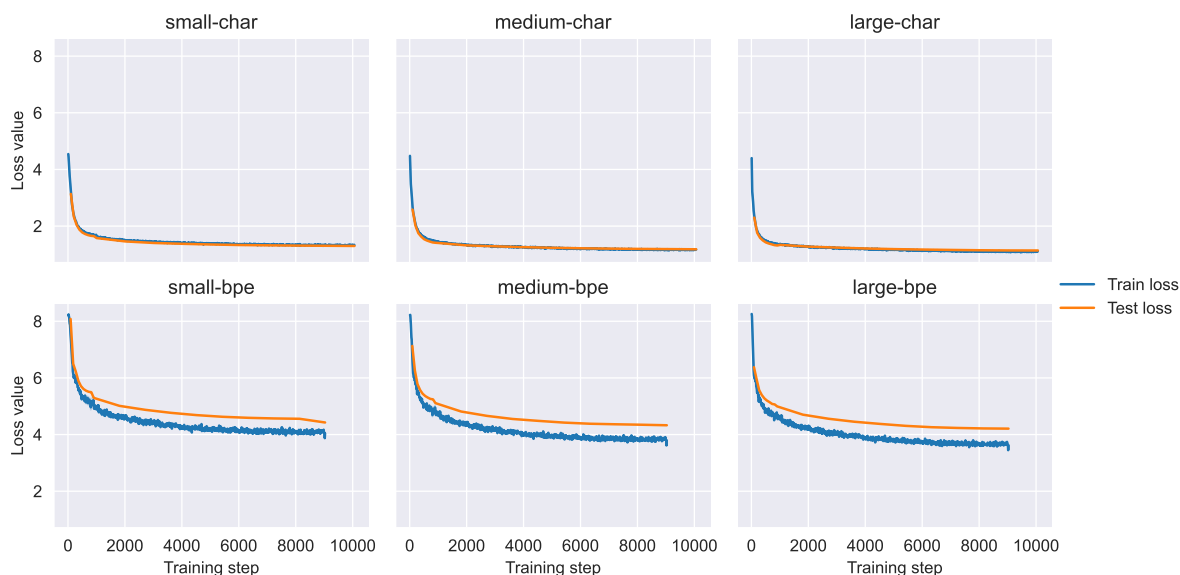


Figure 3: Loss curves for our self-trained Llama models

each individual variable, the other plots are scatterplots which visualize the relationship between the variables. The data points are colored for their tokenization scheme.

In the upper left and lower right plot we can see that both real and non-words are split similarly in the two distinct tokenization schemes. Words split by the BPE tokenizer tend to have fewer tokens, mostly between one and six. For the character-based tokenizer, a normal distribution is visible, with its peak at six tokens.

The upper right and lower left scatter plots show the relationship between tokenization for real and non-words. The character-level tokenization exhibits perfect alignment between both kinds of stimuli, they are always split into the exact same number of tokens. Subword tokenization is slightly skewed towards the non-word tokens. This means that non-words are more often split into more tokens than real words, although the reverse case is not completely infrequent.

C Full learning curves for BLiMP and word learning

Figure 5 shows the full learning curves of all phenomena included in BLiMP (individual syntactic paradigms belonging to one phenomenon are displayed in the same sub-figure) as well as our own lexical benchmarks (all displayed in individual sub-figures), for our six self-trained models and the six Pythia models that we compare them to. We fit a fifth-order polynomial curve to the individual data points and display it on a logarithmic scale.

It should be noted that we plot the number of the checkpoint on the x-scale. However, the individual amounts of actual textual data seen between these checkpoints differs vastly between our self-trained models (10M lexical tokens) and the Pythia models (825GB of textual data; as the dataset has since been taken down, no lexical token counts are possible anymore).

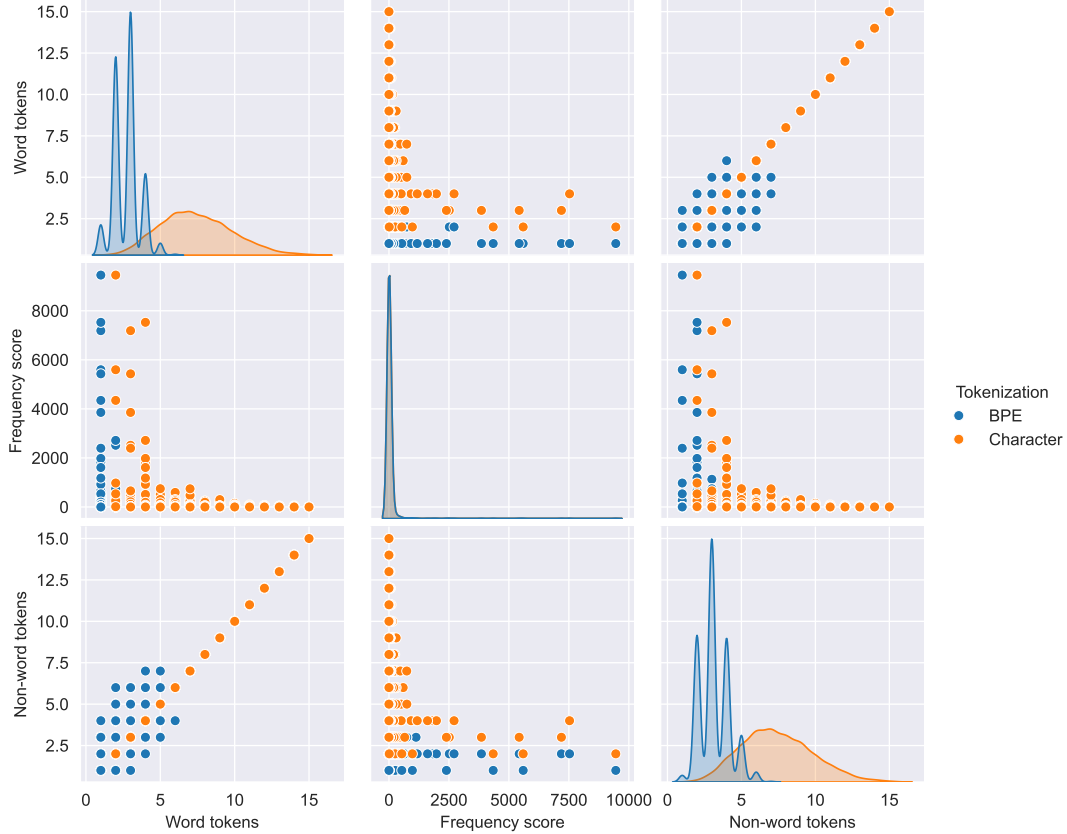


Figure 4: Pairplot displaying (i) number of tokens of words, (ii) frequency scores from CELEX (Baayen et al., 1995) and (iii) number of tokens of non-words (for BPE and character tokenization)

D Correlations between word learning and syntactic learning

As an additional measure of commonalities between word learning and syntactic learning, we calculate Spearman-rank correlation scores between ordered accuracy scores for our lexical tasks and BLiMP paradigms. Table 3 shows the underlying numerical values for the correlation heatmap provided in Figure 6 (please note that the heatmap is rotated in comparison to the table). All scores are statistically significant ($p < 0.05$). Due to the similar learning curves found in Figure 2, we average accuracy scores over all lexical phenomena (lexical decision and both surprisal settings), and then calculate correlations between them (both high and low frequency) and the coarse-grained BLiMP phenomena. For BPE models, lexical performance is highly correlated with more than half of the BLiMP phenomena. The character models show much weaker correlation with syntactic learning. This further confirms our findings about the strong entanglement of lexical and syntactic learning in subword models and their weaker ties

in character models.

E Final BLiMP scores for all models

We reproduce the final syntactic evaluation scores for all models that we incorporated in our lexical analyses in Table 4. Generally, scores improve with larger models and with more training data. Most strikingly, subword models are consistently superior to comparable character models trained on the same amount of data. These differences, however, are most pronounced for the small models trained on very little data, like our Llama models trained on 10M tokens (7% for smallest models, 2% for largest models). For the comparable GPT-2 models trained on 100M tokens, the gap becomes much smaller (0.4%).

F Development of word/non-word differences

In Figure 7, we plot the average difference between word and non-word negative log-probability values across training, for both high-frequency and low-frequency data. Positive scores indicate preference




Figure 5: Learning curves for all paradigms in BLiMP and high/low frequency lexical decision data, separated for models (rows) and phenomenon sets (columns)




Figure 6: Correlation heatmap

for real words. For the character models, the differences are generally less pronounced and get most extreme at the end of pre-training (where accuracy scores do not change anymore), especially for the lexical decision data, which is already consistent at very early training stages. For the BPE models,

we see that at the beginning they actually prefer non-words in the lexical decision task. Only after the first 10% of training they *begin* to discern words and non-words. While overall tendencies remain the same for both frequency conditions, the absolute differences are generally lower and the

BLiMP phenomenon	small-char		medium-char		large-char		small-bpe		medium-bpe		large-bpe	
	highFrq	lowFrq	highFrq	lowFrq	highFrq	lowFrq	highFrq	lowFrq	highFrq	lowFrq	highFrq	lowFrq
Anaphor agr.	-0.393	-0.435	0.580	0.277	0.332	0.555	0.569	0.559	0.930	0.910	0.772	0.759
Argument structure	0.339	0.589	0.467	0.825	0.545	0.895	0.949	0.959	0.970	0.987	0.979	0.986
Binding	-0.718	-0.629	0.660	0.918	0.653	0.922	0.901	0.889	0.992	0.993	0.979	0.977
Control raising	0.904	0.837	0.909	0.776	0.791	0.892	0.777	0.780	0.974	0.974	0.930	0.936
Det.-noun agr.	0.686	0.870	0.524	0.869	0.521	0.890	0.989	0.989	0.990	0.993	0.994	0.989
Ellipsis	-0.912	-0.766	-0.285	0.209	0.180	0.662	0.857	0.822	0.897	0.868	0.865	0.856
Filler gap	-0.765	-0.586	-0.554	-0.146	-0.200	0.209	-0.715	-0.722	-0.589	-0.602	-0.063	-0.031
Irregular forms	0.612	0.724	0.507	0.856	0.397	0.787	-0.116	-0.098	0.636	0.662	0.816	0.832
Island effects	-0.937	-0.840	-0.862	-0.657	-0.556	-0.214	-0.321	-0.334	0.473	0.418	-0.088	-0.094
NPI licensing	0.535	0.641	0.748	0.782	0.616	0.568	-0.425	-0.407	0.520	0.526	-0.069	-0.049
Quantifiers	-0.476	-0.202	0.059	0.299	0.547	0.848	0.789	0.767	0.779	0.770	0.716	0.695
Subj.-verb agr.	-0.386	-0.400	0.329	0.529	0.351	0.730	0.963	0.961	0.982	0.981	0.967	0.974

Table 3: Spearman-rank correlation scores between ordered accuracy scores for our lexical tasks and BLiMP paradigms

Tok.	Model	Params	BLiMP score
Subword (BPE)	Pythia	14M	65.86%
		70M	73.30%
		160M	77.50%
		410M	81.63%
		1B	82.21%
		1.4B	81.92%
	GPT-2	85M	77.80%
	Llama	2.51M	59.80%
		7.77M	64.55%
		30.03M	64.56%
Character	GPT-2	85M	77.40%
	Llama	0.49M	52.69%
		3.73M	51.07%
		21.94M	62.14%

Table 4: BLiMP scores for all models

differences between the curves are less pronounced in the low-frequency setting.

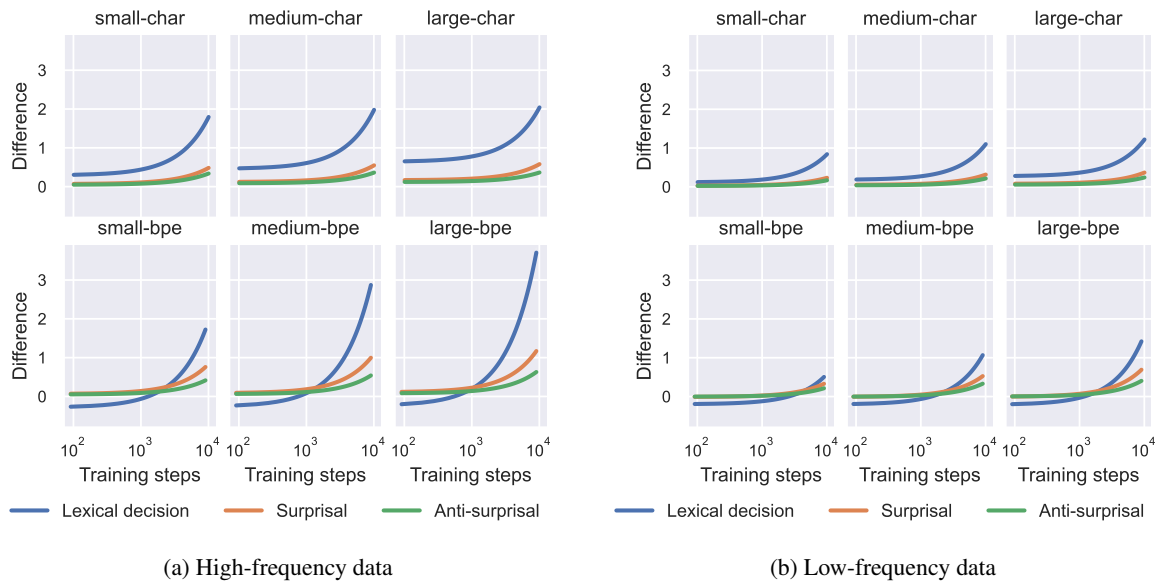


Figure 7: Average differences between surprisal values across pretraining