Is morphological ordering subject to innate constraints? A test of the Poverty of the Stimulus Argument using Transformer-based Language Models

David Silva (fc53755@alunos.ciencias.ulisboa.pt)[1,2], Vânia Mendonça[2], João Veríssimo[1]
[1]Center of Linguistics, School of Arts and Humanities, University of Lisbon, Portugal, [2]LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

The Poverty of the Stimulus Argument holds that children acquire certain linguistic generalisations despite insufficient evidence in the input, suggesting that language acquisition is guided by innate grammatical constraints. Recent work has tested this claim by asking whether Transformer-based Language Models (TLM) trained on developmentally plausible datasets can acquire comparable generalisations. If TLM can capture these generalisations without language-specific innate constraints, this would weaken the argument that such biases are necessary for human language learning [6].

As a test case, we examined the Plurals-in-Compounds (PiC) effect, which refers to English speakers' tendency to avoid regular, but not irregular, plurals inside compounds: for instance, *rats tracker* is less acceptable than *rat tracker*, but this penalty is reduced or absent for *mice tracker* vs. *mouse tracker* [2]. One account holds that this effect reflects an innate constraint on morphological ordering, because it is present in very young children (3–5 year-olds) despite limited evidence in the input [1, 2]. Other accounts claim that the PiC effect can be learned from the distribution of phonological and semantic properties [3, 5].

To test these competing views, we examined whether the PiC effect emerges in TLM trained on a developmentally plausible BabyLM corpus of 10M words, which is comparable to the linguistic input available during the first 1.5 to 3 years of life [6]. If the models exhibit the PiC effect, this would argue against the need for innate constraints on morphological ordering.

Methods: We created 224 novel English noun–noun compounds containing regular and irregular plural non-heads and their singular counterparts (e.g., *rats*/*mice*/*rat*/*mouse tracker*). Non-head nouns were matched between conditions in frequency, length, and semantic properties. All compounds had an unambiguous meaning and were absent from the models' training data.

We used 17 TLM from the BabyLM challenge [6]. Model behavior was evaluated through Surprisal of the head noun (e.g., *tracker* in *rats tracker*): $Surprisal = -\log_2 P(word|context)$. Higher Surprisal was interpreted as lower acceptability. The data were analysed using Bayesian mixed-effects regression, with model as a random factor. Bayes Factors (BF_{10}) were used to quantify the evidence for each effect [4].

Results and Discussion: We obtained evidence for an interaction between regularity and plurality (b = 0.70 [0.40, 0.99], $BF_{10} > 100$). Compounds containing regular plurals showed higher Surprisal than their singular counterparts (b = 0.95 [0.58, 1.31], $BF_{10} > 100$), indicating reduced acceptability (i.e., *rats tracker* was less acceptable than *rat tracker*; Fig. 1a). In contrast, this penalty was much weaker for irregular plurals (b = 0.25 [-0.13, 0.63], BF_{10} = 0.32). The random-effect estimates showed the same general pattern across all 17 models (Fig. 1b).

Our findings indicate that the PiC effect can emerge in TLM trained on developmentally plausible corpora, suggesting that the input contains enough distributional information to support this generalisation without invoking innate constraints.

References

[1] Iris Berent and Steven Pinker. The dislike of regular plurals in compounds: Phonological familiarity or morphological constraint? *The Mental Lexicon*, 2007.

[2] Peter Gordon. Level-ordering in lexical development. *Cognition*, 21, 1985.

[3] Todd R. Haskell, Maryellen C. MacDonald, and Mark S. Seidenberg. Language learning and innateness: Some implications of Compounds Research. *Cognitive Psychology*, 2003.

[4] Michael D. Lee and Eric-Jan Wagenmakers. *Bayesian Cognitive Modeling: A Practical Course*. Cambridge University Press, 1 edition, 2014.

[5] Michael Ramscar and Melody Dye. Learning language from the input: Why innate constraints can't explain noun compounding. *Cognitive Psychology*, 2011.

[6] Ethan G. Wilcox, Michael Y. Hu, Aaron Mueller, Alex Warstadt, Leshem Choshen, Chengxu Zhuang, Adina Williams, Ryan Cotterell, and Tal Linzen. Bigger is not always better: The importance of human-scale language modeling for psycholinguistics. *Journal of Memory and Language*, 2025.

Figures and Tables

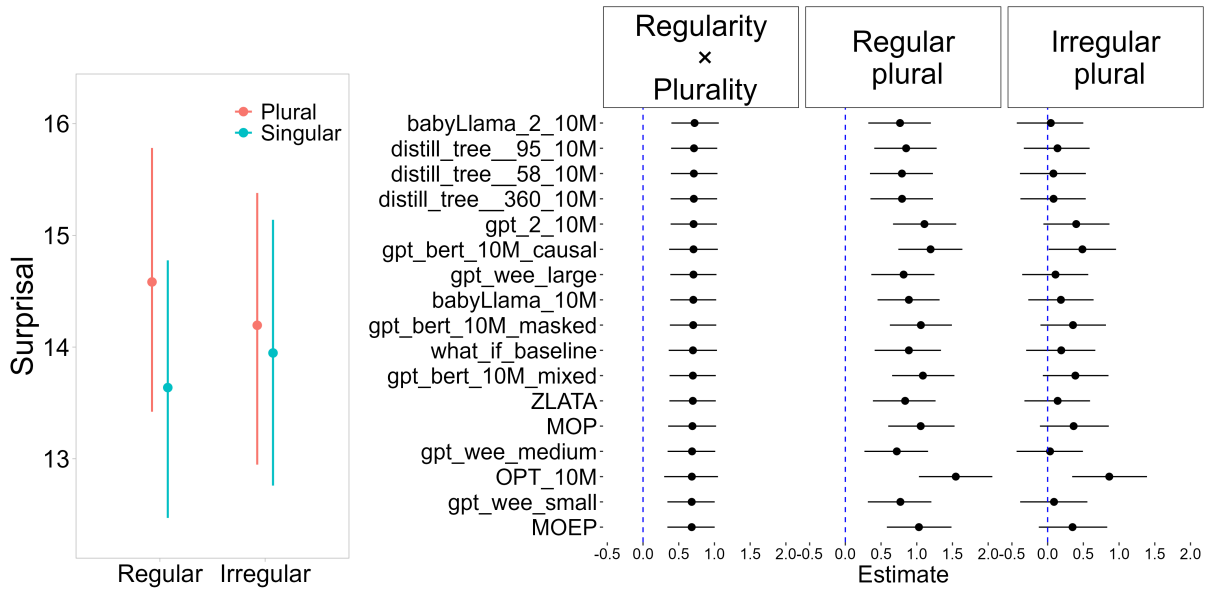

Figure 1: (a) Mean Surprisal values for head nouns across all models; (b) Forest plot with the estimated regular plural effect (e.g., *rat tracker* vs. *rats tracker*), irregular plural effect (e.g., *mouse tracker* vs. *mice tracker*) and regularity × plurality interaction, for each Transformer model.