

Learning from Impairment: Leveraging Insights from Clinical Linguistics in Language Modelling Research

Dominique Brunato

Istituto di Linguistica Computazionale "Antonio Zampolli" (CNR-ILC)

ItaliaNLP Lab, Pisa

dominique.brunato@ilc.cnr.it

Abstract

This position paper investigates the potential of integrating insights from language impairment research and its clinical treatment to develop human-inspired learning strategies and evaluation frameworks for language models (LMs). We inspect the theoretical underpinnings underlying some influential linguistically motivated training approaches derived from neurolinguistics and, particularly, aphasiology, aimed at enhancing the recovery and generalization of linguistic skills in aphasia treatment, with a primary focus on those targeting the syntactic domain. We highlight how these insights can inform the design of rigorous assessments for LMs, specifically in their handling of complex syntactic phenomena, as well as their implications for developing human-like learning strategies, aligning with efforts to create more sustainable and cognitively plausible natural language processing (NLP) models.

1 Introduction

The way human language develops and breaks down offers a privileged lens to investigate central aspects of human language competence, including the nuanced sensitivity to sentence acceptability and complexity (Caramazza and Zurif, 1978; Guijarro-Fuentes, 2020; Crain, 1991; Goodman, 1997). This exploration is not only crucial to test linguistic theories on empirical evidence but also has significant implications for crafting language resources informed by cognitive principles. In the current NLP era dominated by Large Language Models (LLMs), such resources are gaining increasing significance (Opitz et al., 2024), serving as essential benchmarks to unravel the linguistic competence implicitly encoded in neural network representations and possibly shedding light on similarities and differences between how humans and LLMs acquire, represent, and process this knowledge (Warstadt and Bowman, 2022; Belinkov and Glass, 2019; Baroni, 2022). To date,

much research in this direction has been focused on creating targeted diagnostics to assess LLMs' syntactic and semantic abilities inspired by formal linguistics and psycholinguistics (Warstadt et al., 2020, 2019; Ettinger, 2020; Li et al., 2021; Someya et al., 2024). Alternatively, researchers have drawn inspiration from language acquisition research to explore the use of developmental complexity principles in building more efficient language models (Warstadt et al., 2023), or to compare how humans and algorithms learn language, aiming to identify shared patterns (Evanson et al., 2023; Yedetore et al., 2023; Yedetore and Kim, 2024).

Building upon this research paradigm that seeks to investigate how human-like language abilities naturally emerge in neural language models—starting with the creation of well-motivated and rigorous test sets and training methods—this position paper proposes a novel complementary approach. Specifically, we argue for leveraging the realm of speech and language pathology (SLP) and rehabilitation research as rich data sources to gain insights not only for isolating a set of complexity phenomena but also for devising metrics that facilitate the identification of *hierarchies of complexity* within the same phenomenon, with a particular emphasis on syntactic ones. These nuanced complexity metrics may serve dual purposes. On the one hand, they can be used to craft challenging evaluation protocols to assess the proficiency of deep neural networks in handling key factors underlying sentence complexity. On the other one, in line with efforts to develop more sustainable LMs incorporating cognitively plausible mechanisms and trained on reasonably sized data (Treviso et al., 2023; Warstadt et al., 2023; Huebner et al., 2021), they can inform human-inspired language learning strategies, such as those based on curriculum learning (Bengio et al., 2009; Agrawal et al., 2021; Chang et al., 2021), to order training instances based on human-inspired complexity metrics or

to design new, linguistically-informed pre-training tasks or targeted prompts aiming at maximizing the generalization of linguistic skills in models.

2 Linguistic Complexity from the SLP Perspective

Linguistic complexity plays a crucial role in human sentence processing and also holds strong connections to how both native speakers and deep neural network models perceive the grammaticality of a sentence (Linzen et al., 2016; Marvin and Linzen, 2018; Gulordava et al., 2018).

While the long-standing interest for this construct across many communities, defining linguistic complexity is still challenging given its multidimensional and multifarious nature (Brunato et al., 2018). Common distinctions include ‘objective vs. agent-related’ (Dahl, 2004) or ‘absolute vs. relative’ (Miestamo, 2008) complexity. Despite different formulations, the first term of these distinctions considers the formal properties of the linguistic system, whereas the second one addresses the issues of cost, difficulty, and level of demand for a language user/learner. From a developmental perspective, it is assumed that “a more complex item is expected to be acquired later than a less complex one” (Di Domenico, 2017) or, in different words, that linguistic complexity mirrors “the order in which linguistic structures emerge and are mastered in second (and, possibly, first) language acquisition” (Pallotti, 2015). Complexity, in this view, is understood as how an item emerges along stages of development or the time required for it to be processed in an target-like (i.e. adult or native speaker) manner (Friedmann and Reznick, 2021).

Like the developmental perspective, the Speech and Language Pathology perspective is user-based but intersects with the specific challenges individuals face due to developmental or acquired language disorders. These challenges can manifest in difficulties with both language comprehension and production, impacting specific or multiple language domains (Blackwell and Bates, 1995). Within this broad field, our investigation focuses on aphasiology, specifically agrammatism literature and its clinical implications. Agrammatism is a form of Broca’s aphasia, a general linguistic impairment typically resulting from brain damage such as trauma or stroke. It is characterized by structurally impoverished, ‘telegraphic’ speech, with short and fragmented sentences, consisting mainly

of content words, while function words and morphemes are frequently omitted. Comprehension of grammatically complex sentences can also be impaired, though understanding of single words and simple sentences is often intact (Goodglass and Menn, 1985). While the key factors shaping the behavioural patterns of loss and preservation in agrammatic aphasia are still debated, it has been suggested that linguistic complexity plays a crucial role in explaining the deficits that these patients have both in production and comprehension (Bastiaanse et al., 2009; Avrutin, 2001), as well as in maximizing treatment gains and generalization patterns in language recovery.

In line with the purpose of this paper, we intend to highlight how the evidence and theoretical underpinnings underlying some influential treatment protocols grounded on linguistic theory can help uncover cognitively-informed *hierarchies of complexity* in language phenomena that can benefit language modelling research in several ways.

3 Linguistically-Informed Approaches to Language Recovery in Aphasia

Developing language treatments that not only enhance trained items but also promote generalization to untrained items/tasks is a key concern in aphasia research. Generalization, defined and applied in various ways over the last half-century, presents challenges for speech-language pathologists in planning, implementing, and measuring treatment protocols (Mayer et al., 2024). Rather than exhaustively reviewing all rehabilitation methods and their efficacy, which is beyond the scope of this paper¹, this section delves deeper on some frameworks rooted in theoretical linguistics (Garraffa and Fyndanis, 2020), focusing particularly on two of them where linguistic stimuli used in training are grounded in linguistic constructs arranged according to an operationalization of complexity².

3.1 Mapping Therapy (MT)

The MT approach moves from the assumption that sentence production and comprehension impairments in agrammatism are due to difficulties in mapping thematic roles (agent, theme) onto their grammatical constituents (subject, object). This

¹Interested readers can be referred to (Webster et al., 2015; Mayer et al., 2024; Fontoura et al., 2012), *inter alia*.

²Whenever available, all examples are derived from the original stimuli contained in the cited reference papers and presented in Appendix.

can explain the higher failure that these patients exhibit in comprehending and/or producing **non-canonical sentences**, such as passives and object relative clauses, where the mapping process is indirect compared to active sentences, and particularly in sentences with reversible semantic roles. Thus, according to the so-called “mapping hypothesis” (Schwartz et al., 1994), the proposed treatment program is directed at the remediation of these mapping operations through a metalinguistic training task in which patients are trained to recognize the thematic roles and their syntactic realizations in different structures. The mapping therapy approach has inspired numerous targeted training methods, all emphasizing the essential role of verb argument structure and prompting participants to recognize thematic role information for correct sentence production. Generally, these protocols follow a step-wise approach with slight variations, organizing training into stages that correspond to increasing levels of complexity in the target stimuli. The treatment protocol introduced by (Rochon et al., 2005) includes four levels (see Appendix A for examples). In Level 1 participants are presented with pictures as eliciting material and they are trained to recognize and use only the agent cue to produce active sentences, specifically declaratives and subject clefts. In Level 2, the focus shifts to recognizing and using the theme cue to produce passives and object cleft sentences. As a preparatory step to the final level of treatment, Level 3 introduces the identification of both roles, but in a fixed order. Here, sentences can start with either role, allowing for the elicitation of all four sentence structures. In Level 4, role identification is varied and randomized, with either an agent or theme cue provided.

As in most protocols, learning and generalization effects are usually evaluated across stimuli (using the same structures but different input verbs) and/or across tasks.

3.2 Treatment of Underlying Forms (TUF)

TUF is a specific linguistic treatment introduced in (Thompson and Shapiro, 2005) that focuses on noncanonical sentence structures and incorporates training for both sentence production and comprehension. It is based on the premise that training underlying, abstract properties of language facilitates generalization to untrained structures with similar linguistic properties. Like MT, it recognizes the key role of verbs and focused training in sentence argument structure. However, it places

additional emphasis on the syntactic movement operations (framed in Generative Grammar principles) involved in creating grammatically correct noncanonical sentences. The proposed treatment uses the active form of the expected target sentences to train participants at incremental steps: i.) understand and produce the verb and its arguments in each sentence; ii.) move the proper sentence constituents to form target sentence structures; iii.) produce the surface form of the target sentence; iv.) comprehend and produce the verbs and verb arguments in their noncanonical position. What sets TUF apart from other protocols based on explicit learning is its proposal that training more complex structures first yields better outcomes than the traditional method of gradually increasing item complexity. This approach has been systematized in the **Complexity Account of Treatment Efficacy (CATE)**, according to which “training complex structures results in generalization to less complex structures when untreated structures encompass processes relevant to (i.e., are in a subset relation to treated ones)” (Thompson et al., 2003). As mentioned, the construct of syntactic complexity relates to sentences containing a type of unbounded dependencies, i.e. constructions whose correct understanding/processing requires the computation of grammatical relationship between phrases that are pronounced in a position different from the one where they are interpreted. These constructions, also commonly named “filler-gap” dependencies in psycholinguistics, are known to pose significant challenges for human language processing, including in non-impaired speakers. The challenge is particularly pronounced when the “filler” (i.e., the pronounced element) is far from the “gap” (the position in which the element is interpreted). CATE specifically targets noncanonical sentences derived from *wh-movement*, particularly focusing on **object-extracted dependencies** in three forms: object *wh*-questions (WH), object clefts (OC), and object relative clauses (OR) (see Appendix B for examples). It assumes that a hierarchy of increasing complexity can be traced from the first to last structure, which can be explained according to movement operations and resulting parse tree depth (for details, see the reference papers).

The validity of this hierarchy is supported by the patterns of loss observed in agrammatic comprehension, especially in the effectiveness of predicted recovery trajectories. It has been observed that treatment targeting object relative clauses suc-

cessfully generalized to untreated object clefts and object wh-questions, but not the other way around. For the purpose of this paper, we do not aim to argue the clinical effectiveness of CATE compared to other treatments that follow a more intuitive approach, i.e., progressing from simpler to more complex training examples. Instead, we aim to use the findings within this framework as evidence for the existence of syntactic structures that can be systematically ordered by complexity.

3.3 Syntax Stimulation Program (SSP)

This method, also known as Helm Elicited Language Program for Syntax Stimulation (HELPSS) (Helm-Estabrooks, 1981), aims to enlarge the repertoire of grammatical structures by agrammatic aphasic patients. Differently from the previously described protocols, it focuses on practicing the surface form of target structure rather than training the underlying structure and computations. The training protocol is sequentially arranged from repetition to spontaneous production in context. The program targets eleven types of sentences, which are ordered for level of difficulty as follows: Imperative Intransitive; Imperative Transitive; Wh-Interrogative; Declarative Transitive; Declarative Intransitive; Comparative; Passives; Yes-No Questions; Direct-Indirect Object; Embedded Sentences; Future. Each typology is trained within a story completion format comprising two levels of difficulty: the first uses delayed repetition of the expected target phrase, the second removes the benefit of repetition from the patient.

4 Insights from Aphasia Treatment for Language Modelling Research

Our proposal suggests that training protocols for aphasia treatment can play a crucial role in developing more robust and cognitively plausible LMs informed by human language (re-)learning. Specifically, we outline contributions in two main areas covering both training and evaluation frameworks.

4.1 Evaluation via Complexity Hierarchies

So far, a large variety of benchmarks informed by theoretical and experimental linguistics has been introduced to test the linguistic competencies of deep neural language models (Warstadt et al., 2019, 2020; Hu et al., 2020). However, none has tested the model sensitivity to the same phenomenon (e.g. Wh-movement) arranging stimuli at incremental levels of complexity, as proposed e.g. by

the CATE approach. This involves not only manipulating for grammaticality sentences hierarchically arranged at different complexity levels, and testing the model’s response, but also assessing whether these sentences influence the model’s ability to perform high-level reasoning tasks, such as natural language inference.

4.2 Learning Strategies

Curriculum Learning with Human-Inspired Complexity: Ranking training data based on complexity hierarchies derived from neurolinguistics and clinical treatment could be a viable strategy to optimize the curriculum for LMs, mimicking how humans learn progressively from simpler to more complex structures. Such hierarchies can also serve as benchmarks for evaluating the effectiveness of other training data optimization methods, including those inspired by language acquisition scenarios. While studies so far have reported mixed or negative findings on developmental-inspired learning strategies for LM training, the prevailing assumption is that child-directed speech provides an optimal foundation to foster linguistic generalization capabilities in smaller-scale training regimes (Huebner et al., 2021; Martinez et al., 2023; Eldan and Li, 2023). However, a SLP-informed curriculum presents an alternative, data-efficient pretraining technique by systematically organizing training data around explicit syntactic complexity metrics. Unlike the variability inherent in child-directed speech, this approach creates a structured pathway through progressively complex syntactic forms, which could accelerate generalization and adaptability to novel linguistic constructions.

Linguistic Learning Objectives: Aphasia treatment protocols can inform the creation of new pre-training objectives for LMs. These objectives, recast based on linguistic tasks proposed by clinicians, can supplement traditional language modelling approaches by injecting a structured linguistic bias at a foundational level. Notably, while linguistically-enhanced pre-training tasks like semantic role labeling have shown promise (Cui et al., 2022; Zhou et al., 2020), their effectiveness also combined with training on progressively more complex sentences in terms of diverse argument structures, as in the MT protocol, remains unexplored.

Targeted Prompts for Improved Generalization: Insights from aphasia research can guide the creation of targeted prompts for use in knowledge distillation scenarios. These prompts, designed to

address specific linguistic constructs, improve the model’s capability to generalize and master complex structures. Similar to aphasia therapy, where patients progress through increasingly complex linguistic tasks to (re-)build language abilities, this approach can be applied in LLMs. Structuring prompts to begin with simpler forms and gradually introduce more complex sentence constructions may enable the student model to effectively learn a wide range of linguistic phenomena.

Limitation and Future Directions

As an opinion paper, our work proposes a framework for language modelling research without providing extensive empirical evidence through comparisons with other training strategies or evaluation settings for language models. We have not discussed possible implementation strategies in detail, but potential approaches could include starting from available examples in the literature to identify relevant templates and using synthetic data inspired by them to augment data. Alternatively, sentence constructions such as those listed in the SSP approach might also be searched for in available treebanks for different languages. Notably, while we focus on examples from the original papers (often in English), adaptations exist for many other languages. Moreover, our primary focus was on linguistically-based training methods developed in the field of aphasia, which involve explicit teaching of grammatical structures or syntactic rules to promote generalization and foster metalinguistic knowledge primarily through textual input, often supported by picture only. However, these approaches, while significant, are not unique within the field of aphasia rehabilitation and can be enhanced by incorporating other modalities input alongside the linguistic one. Integrating such multimodal training objectives into language models is a promising avenue for future research.

Acknowledgments

The author gratefully acknowledges the support of the project “LODE – The Language Of Dreams: the relationship between sleep mentation, neurophysiology, and neurological disorders ” - PRIN 2022 2022BNE97C_SH4_PRIN2022 funded by the Italian Ministry of University and Research, and of the project XAI-CARE - PNRR-MAD-2022-12376692 under the NRRP MUR program funded by the NextGenerationEU.

References

- Ameeta Agrawal, Suresh Singh, Lauren Schneider, and Michael Samuels. 2021. [On the role of corpus ordering in language modeling](#). In *SUSTAINLP*.
- Sergey Avrutin. 2001. Linguistics and agrammatism. *Glott international*, 5(3):1–11.
- Marco Baroni. 2022. On the proper role of linguistically oriented deep net analysis in linguistic theorising. In *Algebraic structures in natural language*, pages 1–16. CRC Press.
- Roelien Bastiaanse, Gosse Bouma, and Wendy Post. 2009. [Linguistic complexity and frequency in agrammatic speech production](#). *Brain and Language*, 109(1):18–28.
- Yonatan Belinkov and James Glass. 2019. [Analysis Methods in Neural Language Processing: A Survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Arshavir Blackwell and Elizabeth Bates. 1995. Inducing agrammatic profiles in normals: Evidence for the selective vulnerability of morphology under cognitive resource limitation. *Journal of Cognitive Neuroscience*, 7(2):228–257.
- Dominique Brunato, Lorenzo De Mattei, Felice Dell’Orletta, Benedetta Iavarone, and Giulia Venturi. 2018. [Is this sentence difficult? do you agree?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2690–2699, Brussels, Belgium. Association for Computational Linguistics.
- Alfonso Caramazza and Edgar B. Zurif. 1978. [Language acquisition and language breakdown : parallels and divergencies](#).
- Ernie Chang, Hui-Syuan Yeh, and Vera Demberg. 2021. Does the order of training samples matter? improving neural data-to-text generation with curriculum learning. *CoRR*, abs/2102.03554.
- Stephen Crain. 1991. Language acquisition in the absence of experience. *Behavioral and Brain Sciences*, 14(4):597–612.
- Yiming Cui, Wanxiang Che, Shijin Wang, and Ting Liu. 2022. [Lert: A linguistically-motivated pre-trained language model](#). *ArXiv*, abs/2211.05344.
- Östen Dahl. 2004. *The Growth and Maintenance of Linguistic Complexity*. John Benjamins Publishing, Amsterdam.
- Elisa Di Domenico. 2017. *Syntactic Complexity from a Language Acquisition Perspective*. Cambridge Scholars Publishing, Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK.

- Ronen Eldan and Yuanzhi Li. 2023. [Tinystories: How small can language models be and still speak coherent english?](#) *CoRR*, abs/2305.07759.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models.](#) *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Linnea Evanson, Yair Lakretz, and Jean-Rémi King. 2023. [Language acquisition: do children and language models follow similar learning stages?](#) In *Annual Meeting of the Association for Computational Linguistics*.
- Denise Ren da Fontoura, Jaqueline de Carvalho Rodrigues, Luciana Behs de Sá Carneiro, Ana Maria Monção, and Jerusa Fumagalli de Salles. 2012. Rehabilitation of language in expressive aphasia: a literature review. *Dementia & Neuropsychologia*, 6:223–235.
- N Friedmann and J. Reznick. 2021. [Stages rather than ages in the acquisition of movement structures: Data from sentence repetition and 27696 spontaneous clauses.](#) *Glossa: a journal of general linguistics*, 39(1):1–69.
- Maria Garraffa and Valantis Fyndanis. 2020. Linguistic theory and aphasia: An overview. *Aphasiology*, 34(8):905–926.
- Harold Goodglass and Lise Menn. 1985. Is agrammatism a unitary phenomenon? In *Agrammatism*, pages 1–26. Elsevier.
- Elizabeth Bates Judith C Goodman. 1997. On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia and real-time processing. *Language and cognitive Processes*, 12(5-6):507–584.
- Pedro Guijarro-Fuentes. 2020. Language acquisition and linguistic theory: When linguistic theory meets empirical data. *Applied Linguistics Review*, 11(3):369–376.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. *CoRR*, abs/1803.11138.
- N. Helm-Estabrooks. 1981. *Helm Elicited Language Program for Syntax Stimulation*. Pro-Ed.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. [A systematic assessment of syntactic generalization in neural language models.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Philip A. Huebner, Elior Sulem, Fisher Cynthia, and Dan Roth. 2021. [BabyBERTa: Learning more grammar with small-scale child-directed language.](#) In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 624–646, Online. Association for Computational Linguistics.
- Bai Li, Zining Zhu, Guillaume Thomas, Yang Xu, and Frank Rudzicz. 2021. [How is BERT surprised? layerwise detection of linguistic anomalies.](#) 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 4215–4228, Online. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies.](#) *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Richard Diehl Martinez, Hope McGovern, Zebulon Goriely, Christopher Davis, Andrew Caines, Paula Buttery, and Lisa Beinborn. 2023. [CLIMB – curriculum learning for infant-inspired model building.](#) In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 112–127, Singapore. Association for Computational Linguistics.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Jessica F. Mayer, Erin B. Madden, Jennifer Mozeiko, et al. 2024. [Generalization in aphasia treatment: A tutorial for speech-language pathologists.](#) *American Journal of Speech-Language Pathology*, 33(1):57–73.
- Matti Miestamo. 2008. Grammatical complexity in a cross-linguistic perspective. In Matti Miestamo, Kaius Sinnemäki, and Fred Karlsson, editors, *Language Complexity: Typology, Contact, Change*. John Benjamins Publishing, Amsterdam.
- Juri Opitz, Shira Wein, and Nathan Schneider. 2024. Natural language processing relies on linguistics. *arXiv preprint arXiv:2405.05966*.
- Gabriele Pallotti. 2015. [A simple view of linguistic complexity.](#) *Second Language Research*, 31(1):117–134.
- Elizabeth Rochon, Laura Laird, Arpita Bose, and Joanne Scofield. 2005. Mapping therapy for sentence production impairments in nonfluent aphasia. *Neuropsychological rehabilitation*, 15(1):1–36.
- Myrna F Schwartz, Eleanor M Saffran, Ruth B Fink, Jessica L Myers, and Nadine Martin. 1994. Mapping therapy: A treatment programme for agrammatism. *Aphasiology*, 8(1):19–54.
- Taiga Someya, Ryo Yoshida, and Yohei Oseki. 2024. [Targeted syntactic evaluation on the Chomsky hierarchy.](#) In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language*

- Resources and Evaluation (LREC-COLING 2024)*, pages 15595–15605, Torino, Italia. ELRA and ICCL.
- Cynthia Thompson and Lewis Shapiro. 2005. Treating agrammatic aphasia within a linguistic framework: Treatment of underlying forms. *Aphasiology*, 19(10-11):1021–1036.
- Cynthia K Thompson, Lewis P Shapiro, Swathi Kiran, and Jana Sobeks. 2003. [The role of syntactic complexity in treatment of sentence deficits in agrammatic aphasia: the complexity account of treatment efficacy \(cate\)](#). *J Speech Lang Hear Res.*, 46(3):591–607.
- Marcos Treviso, Ji-Ung Lee, Tianchu Ji, Betty van Aken, Qingqing Cao, Manuel R. Ciosici, Michael Hassid, Kenneth Heafield, Sara Hooker, Colin Raffel, Pedro H. Martins, André F. T. Martins, Jessica Zosa Forde, Peter Milder, Edwin Simpson, Noam Slonim, Jesse Dodge, Emma Strubell, Niranjan Balasubramanian, Leon Derczynski, Iryna Gurevych, and Roy Schwartz. 2023. [Efficient Methods for Natural Language Processing: A Survey](#). *Transactions of the Association for Computational Linguistics*, 11:826–860.
- Alex Warstadt and Samuel R. Bowman. 2022. [What artificial neural networks can tell us about human language acquisition](#). *Preprint*, arXiv:2208.07998.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan 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 1–34, Singapore. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohanane, 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.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- J Webster, Anne Whitworth, and J Morris. 2015. Is it time to stop “fishing”? a review of generalisation following aphasia intervention. *Aphasiology*, 29(11):1240–1264.
- Aditya Yedetore and Najoung Kim. 2024. [Semantic training signals promote hierarchical syntactic generalization in transformers](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4059–4073, Miami, Florida, USA. Association for Computational Linguistics.
- Aditya Yedetore, 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](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9370–9393, Toronto, Canada. Association for Computational Linguistics.
- Junru Zhou, Zhuosheng Zhang, Hai Zhao, and Shuailiang Zhang. 2020. [LIMIT-BERT : Linguistics informed multi-task BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4450–4461, Online. Association for Computational Linguistics.

A Examples from the Mapping Therapy Protocol (Rochon et al., 2005)

Level 1, Agent Role Cue: Target sentences are reversible, 2-argument structures in the form of actives and subject clefts.

Target sentence: “The nurse chases the tall teacher”

Examiner says:

- This is a picture about chasing.
- The verb in the sentence is ‘chases’.
- In this picture the one being chased is the tall teacher (Theme).
- The one doing the chasing is the nurse (Agent)
- Please make a sentence starting with the nurse*.

Level 2, Theme Role Cue: Target sentences are reversible, 2-argument structures in the form of passives and object clefts.

Target sentence: “The farmer was hugged by the soldier”.

Examiner says:

- This is a picture about hugging.
- The verb in the sentence is “hugged”.
- In this picture the one being hugged is the farmer.
- The one doing the hugging is the soldier.
- Please make a sentence starting with the farmer.

Level 3, Agent and Theme Role Cue: Target sentences are reversible, 2-argument structures in all of the four possible forms: actives, subject clefts, passives, object clefts.

Target sentence: “The baker was called by the judge”.

Examiner says:

- This is a picture about calling.
- The verb in the sentence is “called”.
- In this picture the one doing the calling is the judge.
- The one being called is the baker.
- Please make a sentence starting with the baker.

Level 4, Agent or Theme Role Cue: Target sentences are reversible, 2-argument structures in all the four possible forms: actives, subject clefts, passives, object clefts.

Target sentence: “It is the author that is shooting the farmer”.

Examiner says:

- This is a picture about shooting.
- The verb in the sentence is “shooting”.
- In this picture the one doing the shooting is the author.
- The other person is the farmer.
- Please make a sentence starting with the author.
- Please start your sentence with “It is”.

* At all levels, when the target was a subject cleft or object cleft sentence, the examiner added: “Please begin the sentence starting with ‘it is’.”

B Examples from the Complexity Account of Treatment Efficacy (CATE) (Thompson et al., 2003)

- (1) a. (OR) The man saw the thief who the artist chased _.
- b. (OC) It was the thief who the artist chased _.
- c. (WH) Who did the artist chase _?
- (2) a. (OR) The man saw the husband who the wife covered _.

- b. (OC) It was the husband who the wife covered _.
 - c. (WH) Who did the wife cover _?
- (3)
- a. (OR) The man saw the guest who the waiter watched _.
 - b. (OC) It was the guest who the waiter watched. _.
 - c. (WH) Who did the waiter watch_?