

Is it the end of (generative) linguistics as we know it?

Cristiano Chesi

NeTS-IUSS Pavia, Italy <cristiano.chesi@iusspavia.it>

A significant debate has emerged in response to a paper written by Steven Piantadosi (Piantadosi 2023) and uploaded to the LingBuzz platform, the open archive for generative linguistics. Piantadosi's dismissal of Chomsky's approach is ruthless, but generative linguists deserve it. In this paper, I will adopt three idealized perspectives – computational, theoretical, and experimental – to focus on two fundamental issues that lend partial support to Piantadosi's critique: (a) the evidence challenging the Poverty of Stimulus (PoS) hypothesis and (b) the notion of simplicity as conceived within mainstream Minimalism. In conclusion, I argue that, to reclaim a central role in language studies, generative linguistics – especially Minimalism, which represents a prototypical theoretical perspective on language – needs a serious update leading to (i) more precise, consistent, and complete formalizations of foundational intuitions and (ii) the establishment and utilization of a standardized dataset of crucial empirical evidence to evaluate the theory's adequacy. On the other hand, ignoring the formal perspective leads to major drawbacks in both computational and experimental approaches. Neither descriptive nor explanatory adequacy can be easily achieved without the precise formulation of general principles that can be challenged empirically.

KEYWORDS: Minimalist Grammars, Poverty of Stimulus, Merge, T-model, Descriptive Adequacy, Explanatory Adequacy, Simplicity, Minimum Description Length.

1. Introduction

In November 2013, a group of prominent formal linguists¹ was invited by Andrea Moro to the University School for Advanced Studies IUSS Pavia to discuss foundational issues of generative grammar. The bombastic title of the roundtable I helped organize was 'Can there be a Hilbert List of Syntax (yet)?'. A few months before, Noam Chomsky delivered an inspiring opening lecture for the 2012 academic year at IUSS on 'Language and the Limits of Understanding'. Stimulated by his insightful commentary, we approached the roundtable with considerable excitement, ready to delve into the current state of generative linguistics and to outline its history of success. Just to mention three highly influential ideas whose proposers were present at that roundtable, we could list: (a) The featural sensitivity of locality, known as Relativized Minimality (Rizzi 1990); a very pervasive and persuasive restriction on

specific kinds of non-local dependencies with a clear impact not only on the elegance of the grammatical competence description, but also on certain performance issues – e.g. under the insightful interpretation of relativized minimality in terms of featural inclusion (Starke 2001), the difficulty in processing object relative clauses can be attributed to the challenge of differentiating distinctive features (Friedmann *et al.* 2009; Grillo 2008). (b) The universal ordering of functional categories (Cinque 1999; Rizzi 1997), despite their apparent cross-linguistic variation; a study that originated a fruitful inquiry program dubbed Cartography (Belletti 2004; Cinque 2002; Rizzi 2004; Rizzi & Cinque 2016 a.o.). (c) The complicated – but stable – relationship between hierarchy and linear order, as defined by Antisymmetry (Kayne 1994); a formal intuition that restricts the number of plausible structural descriptions logically admitted on the basis of a dominance-precedence mapping concern. If the number of options for pairing hierarchical structure with the linearly ordered set of pronounced (or signed) morphemes – necessarily imposed by the sensory-motor system – were restricted, the ‘logical problem of language acquisition’ would be more manageable and potentially more solvable under the Principles and Parameters perspective (Chomsky 1981).

At the end of two days of discussion, after 16 brilliant talks, we tried to take stock of the proposed problems agenda, but our attempt was rather disappointing: although each problem was apparently crucial and well stated, the extension of the relevant empirical basis fitting the specific theory set-up was sometimes hard or faint. More crucially, it was practically impossible to present all the problems in a concise and coherent manner within a consistent framework: in nearly every instance, although each problem statement came with a proposed solution, the underlying assumptions were often at odds with the premises of others.

A spectre was haunting generative linguistics – the spectre of Minimalism. In fact, this turned out to be a sum of idiosyncratic interpretations.

In Chomsky’s intuition, Occam’s razor was needed to purge the evolution of Government and Binding Theory (Chomsky 1981) of unnecessary machinery. As a result, the theory should have been simpler, mathematically sound, computationally efficient, and include nothing but what is strictly necessary. This approach marked the beginning of a new research program dubbed Minimalism (Chomsky 1995). Thirty years on, it must be acknowledged that while the program began with commendable intentions, the emerging framework still lacks consistency, especially

regarding the application of fundamental structure-building operations to empirical problems.

In the end, events have surpassed the intentions of the roundtable. A thorough investigation into the success of very Large Language Models (vLLMs),² alongside statistical and experimental advancements, may lead one to agree with many in concluding that generative linguistics no longer dictates the agenda for future linguistic challenges. This position is summarized by Steven Piantadosi who fosters the idea that vLLMs are “genuine theories of language, including representations of syntactic and semantic structure” (Piantadosi 2023). Although these models are primarily designed for a wide range of Natural Language Processing tasks, from Machine Translation to Question Answering, Piantadosi’s key argument is their superiority as linguistic theories. According to him, these models surpass generative approaches in performing comprehensive syntactic tests, ranking vLLMs as the most effective linguistic theories currently available. One platform designed for performing such linguistic benchmarks is SyntaxGym (Hu *et al.* 2020): an on-line, open-source repository that includes a significant set of linguistic contrasts – 39 test suites that include a total of about 4k sentences. For each relevant contrast included, human generalizations have been gathered in various studies. Direct comparisons of these data with the predictions provided by the models under evaluation is then possible. An exemplificative contrast included in SyntaxGym involves non-local agreement dependency, where the subject must agree in number with the matrix copula, despite being linearly separated by a relative clause (from Marvin & Linzen 2018):

(1) a. *The author* that the senators hurt *is* good
b. **The author* that the senators hurt *are* good

A model predicting that (1a) ‘is better than’ (1b) represents a more adequate theory with respect to a theory that is not able to infer the ungrammaticality of (1b). The utility of these minimal pairs in linguistic theorizing is uncontroversial and can be further sophisticated. In the example below, for instance, another kind of restriction on non-local dependencies is considered, also known as Across-The-Board extraction, ATB (Williams 1977):

(2) a. I know *what* the guy broke $_\!$ accidentally and the mechanic fixed $_\!$ skilfully.
b. **I know what* the guy broke $_\!$ accidentally and the mechanic fixed *the engine* skilfully.

This ATB constraint predicts that when a *wh*- item ('what') is extracted from the first conjunct ('the guy broke _ accidentally'), a gap coindexed with the same *wh*- item should be present also in the second conjunct ('the mechanic fixed _ skilfully'). The relevant contrast, again included in SyntaxGym, compares the correct configuration (2a) with an ungrammatical minimal variation in which the second gap is filled by an intrusive argument ('the engine'), in (2b).

Although the Minimalism framework is explicitly designed to address these issues, paradoxically, I concur with Piantadosi's critique regarding its inability (at least in mainstream set-ups) to perform adequately in similarly complete and extensive benchmarks as the ones presented in SyntaxGym.³

To frame this problem, we need to consider three major perspectives, albeit somewhat idealized for the sake of discussion: first, the computational perspective, which posits that the best linguistic theory is simply the one that performs optimally on a shared test set. Second, the theoretical perspective, that considers this goal (i.e. observational adequacy) just as a starting point, being the final goals descriptive adequacy – the theory should be grounded in robust 'genuine generalizations' (Chomsky 2021a) – and explanatory adequacy – the theory should account for language learnability. Lastly, the experimental perspective, that underscores the imperative of meticulous data collection and analysis, signaling a departure from purely theoretical 'armchair linguistics'. I will contend that, unless integrated, these perspectives individually lead to impasses.

In fact, all these three perspectives will be essential to reexplore two foundational issues (section §3): the Poverty of Stimulus (PoS) hypothesis (§3.1) and the notion of simplicity applied to structure-building operations (Merge and Move) in mainstream Minimalism (§3.2). In both cases, I will provide logical and empirical evidence suggesting that the classic arguments must be re-worked, and the three perspectives are all necessary to avoid confounds.

My primary concern is that the Minimalist Program's underspecification of key concepts, including simple but effective restrictions on the application of structure-building operations, has become untenable (formalization issue). Furthermore, the general underevaluation of experimental and computational advancements by leading scholars in the generative field has contributed to a perception of generative linguistics as marginal within both computational and experimental language research communities. I will argue here that it is imperative to bridge the formalization gaps as effectively as possible and to adopt a modern approach to theory evaluation that relies on shared datasets and metrics

(evaluation issue). On the other hand, what effectively guides sound (linguistic) inquiry is the search for empirical evidence that rejects a specific theoretical setup, while avoiding confirmation bias as much as possible. In this respect, generative linguists benefit from both valuable experience and clear principles which, once explicitly formalized, can be disproved. In my opinion, this remains the only effective way to build descriptively adequate theories.

To address these issues, the next section (§2) will set the stage by defining the core minimalist concepts and the empirical data that modern methods have made available. I will first introduce the ‘Language Problem’ that any theory must confront. Subsequently, I will discuss the widely adopted ‘T-model’ (Chomsky 1981; Chomsky, Seely *et al.* 2023) and examine the criteria for achieving descriptive adequacy and computational efficiency. I will then explore a four-way classification of relevant empirical data (§2.2.1-§2.2.4). This will pave the way towards a re-analysis of the PoS argument, emphasizing that the only relevant source of linguistic information for explaining language acquisition, and thereby targeting explanatory adequacy, remains positive linguistic evidence found in child-directed speech. While assessing a theory’s descriptive adequacy requires all possible experimental evidence, from the learnability perspective, any implicit metalinguistic information is deemed irrelevant.

2. Empirical Evidence for a Theoretical Perspective

2.1. Setting the Stage: From Descriptive Adequacy to Efficiency Considerations

The primary goal of any linguistic theory (X) is to precisely circumscribe the infinite set (language L) composed by those sentences (Ss) judged as grammatical by native speakers. By ‘sentence’ we simply refer to a (compositionally) interpretable and producible (through signs or sounds) ordered string⁴ of words/morphemes. On this basis, we can define the Language Problem:

Definition 1. Language Problem

Is theory X capable of generating and recognizing all and only the sentences Ss belonging to language L?

A Minimalist framework is a theoretical perspective promoting the shortest possible list of ‘instructions’⁵ that would enable a grammatical

theory to solve the Language Problem, for any natural language L . A theory that solves the Language Problem is considered ‘observationally adequate’ (Chomsky 1964).

In structural terms – that is, an abstract and explicit description of the generalizations based on the observed Ss in L –, a Minimalist Grammar (MG) defines an infinite set of derivations (sequences of steps, D_s) obtained through the applications of essentially one simple structure building operation (*Merge*) over lexical items (l_i) selected from the language lexicon (Lex_L), as exemplified in (3):

(3) Derivation (D_s) of the sentence S : ‘Alice scolds Bill’
i. $Select(Alice, Bill, scolds)$ where $\{Alice, Bill, scolds\} \in Lex_{English}$
ii. $Merge(scolds, Bill) = \{scolds, Bill\}$
iii. $Merge(\{scolds, Bill\}, Alice) = \{Alice, \{scolds, Bill\}\}$

We usually adopt a concise syntactic tree-like structural description, T_s in (4), to represent the history of the derivation D_s presented in (3).

(4) Syntactic representation (T_s) of S as derived in (3)

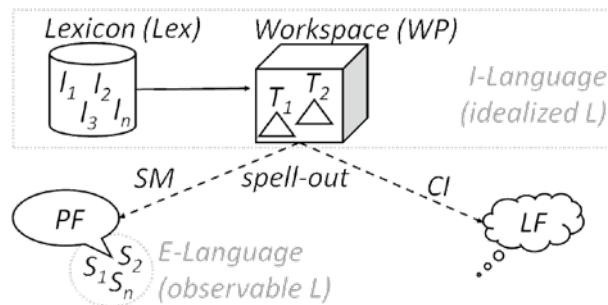


Following the generative grammar tradition, a MG is a theory of I(nternal)-language, modeling the competence of an idealized speaker (Chomsky 1965: 24): this minimally comprises a lexicon specification (Lex_L) and the procedure to obtain appropriate derivations resulting in structural descriptions. Those descriptions should be coherent with the way we pronounce – e.g. /'ælis skəuldz bil/ – and understand – e.g. ‘Alice’ is the agent X, ‘Bill’ is the patient Y, ‘scolds’ the predicate, then ‘X scolds Y’ – the corresponding sentences Ss in L – e.g. ‘Alice scolds Bill’. In this sense, T_s should be a legible scaffolding both for a Phonetic Form, PF, at the Sensory-Motor interface SM, and for interpretation in terms of Logical Form, LF, a compositionally legible structure at the Conceptual-Intentional interface, CI.

Once a pronounceable/interpretable fragment of phrase structure is created – using the structure building operation *Merge* in a workspace, WP –, this fragment is delivered – ‘spelled-out’ – to the two interfaces: CI on the one side and SM on the other. The core syntactic engine will continue to build structures, while the two external modules will independently elaborate on those representations.

These components can be organized under the standard ‘T-’ (or ‘Y-’) model elaborated in (5). Notice that the observable Ss sentences – E(xternal)-language – forming the empirical restriction of the L set are placed at PF, which is the sole stage at which we can empirically observe them. Crucially, the idealized L identified as I-Language, plus externalization constraints (e.g. linear order at PF), should exactly correspond to that observable E-language L.

(5) The ‘T-model’ (re-adapted and extended from Chomsky, Seely *et al.* 2023)



The Minimalist Program was the result of a slimming diet for grammatical theories that appeared to be necessary, at the beginning of the '90s, due to two independent reasons: the first reason was related to the idea of ‘perfection’ (Chomsky 1995: 1), the second to the “proliferation of unwanted derivations” (Chomsky 1995: 283).

As far as the notion of ‘perfection’ is concerned, looking for a theory that only includes the strictly necessary assumptions to satisfy the interface requirements (at SM and CI) is simply restating a long-standing logical problem formulated in many equivalent ways. For instance, in Minimum Description Length, MDL, terms (Grünwald 2007; Rissanen 1978, 1987), or Solomonoff-Kolmogorov-Chaitin complexity (Chaitin 1969; Kolmogorov 1963; Solomonoff 1960), the best theory for a specific empirical domain is the one that better compresses the data observed in that domain. Despite an undecidability menace (M. Li & Vitányi 2008), this is a practical mathematical way to compare the ‘descriptive adequacy’ fit (Chomsky 1965) of a theory (a step forward with respect to ‘observational adequacy’) that subsumes the notion of ‘simplicity’ (Chomsky 2021a):

Definition 2. Descriptive adequacy

the theory that captures more data (S_s in L) with fewer instructions (driving D_s to derive T_s) is the one descriptively more adequate.

The second reason was related to computability: a linguistic theory must be computationally efficient, that is, we should prefer (computable) theories that perform a relevant derivation in the most economical way (Chomsky 1995: 8-9).⁶

Definition 3. Computational efficiency

the computationally most efficient derivation D is the one obtaining the intended structure T in fewer steps.

The distinction between ‘fewer steps’ and ‘fewer instructions’ is critically important: the former refers to the ‘time complexity’ of an algorithm (Barton *et al.* 1987), while the latter pertains to the dimensions of the theory. To understand time complexity, assume a serial procedure where each operation is performed sequentially. Considering a hypothetical scenario in which a very slow processor is capable of executing just one operation per second, and assuming that both Select and Merge count as single operations, executing the procedure (3) would require exactly three seconds. ‘Fewer instructions’, on the other hand, refers to the size of the theory which, in simple terms, can be expressed by the number of bits required to encode the lexicon, the structure-building operations, and the procedure driving the derivations. For example, in Python, less than 300 bytes are needed to encode the instructions for the derivation described in (3).⁷

A generative linguist might feel uncomfortable with these definitions and try to back off under the ‘competence-performance’ shield: the goal of any linguistic theory should be restricted to defining the domain of computation, known as the competence level – or computational level, in Marr’s terminology (Marr 1982) –, and should not encompass the performance level – typically associated with Marr’s algorithmic level. This defense falls short in two ways: first, even though the sequence of operations applied to obtain a certain derivation is assumed to be abstract, with no real temporal implications whatsoever (Chomsky 1995: 380), computational efficiency is included in any MG version – e.g. the ‘phase’ idea⁸ (Chomsky 2001); this implies that algorithmic complexity is at issue here with no necessity to invoke performance. Second, considering the classic representational *vs* derivational opposition⁹ one might still prefer ‘representational filters’ on a complete phrase structure T , once it is fully built – e.g. C-command constraints –, but this approach is logically disadvantaged compared to any derivational constraint that

would exclude the possibility of deriving unwanted structures Ds well before those ill formed descriptions will be later, representationally, valued then discarded to obtain T. This seems to me a straightforward argument on the superiority of derivational approaches over representational ones in computational terms (cf. note 6). One might be suspicious about the mere numeric comparison suggested by Definition 2, but we will see soon that this comparison can be safely restricted to truly recursive generalizations all equally able to generate and recognize an infinite set of data – e.g. relative clause formation.

To summarize, descriptive adequacy (Definition 2) and computational efficiency (Definition 3) pose theoretical boundaries to the Language Problem (Definition 1). While the theoretical perspective tends to idealize the empirical domain to enhance descriptive adequacy (I-language), this approach faces two challenges: firstly, experimental methods tend to enrich the empirical domain in ways that do not always align with idealized contrasts; secondly, the computational perspective, enhanced by machine learning methods, is naturally suited to narrow the gap between idealized language and the observed language L, efficiently and effectively processing any observed sentence S. Let us explore these challenges in more detail in the next section.

2.2. Competence, Performance, and Explanatory Adequacy

After the cognitive revolution (Chomsky 1959), a precise method emerged to approach the study of high cognitive functions. This method was used to avoid imprecise theoretical notions and to obtain fruitful generalizations from the limited set of observable/available data. Under the definitions we just provided in §2.1, one might well imagine that our language faculty not be so perfect – as correctly programmatically assumed in principle by Minimalism. In this case, the computational search for perfection might lead us to solutions to the Language Problem that are more performant than the ones nature has provided.

A classic example is motion: we quickly realized that wheels were more efficient than legs for moving an object from position A to position B. ‘More efficient’ can be interpreted in terms of faster transitions from A to B, more robust – in terms of failures during the motion –, or less complex solutions – simpler design and control. A ‘minimalist program’ for human motion, focusing solely on computational efficiency, would have favored wheels over legs.¹⁰ For this reason, the Language Problem (Definition 1) must be coupled with another fundamental constraint to limit the greediness of Computational Efficiency (Definition 3): we do not want to derive random structure – *mutatis mutandis*: any possible

motion planning from A to B in arbitrary conditions –, but just those structures that are relevant descriptions for the sentences in our language L. From the perspective of the ‘motion problem’, the solution of using wheels would not be suitable for tasks like climbing stairs, which humans, in fact, perform quite well under standard circumstances. At the same time, the temporal predictions of transitions from A to B will be greatly underestimated – wheel-based vehicles would take much less time to move from A to B on a highway than humans.

Even though computational efficiency is, *per se*, a legitimate independent goal, as cognitive scientists we are primarily interested in understanding our language faculty – I-language description –, that is, producing a theory that is, at least, descriptively adequate (Definition 2). Therefore, a trade-off between descriptive adequacy and computational efficiency emerges, possibly favoring a theory that is less compact or efficient but fits better with the data associated with the relevant derivations of sentences Ss in L.

We usually refer to ‘Performance’ as the actual usage of our linguistic knowledge (Chomsky 1965), that is, the way native speakers perform a recognition/generation of sentences Ss in their language L. Ideally, this performance should be precisely predicted by applying externalization considerations – e.g. linear order on structured lexical items – and specific resources constraints – e.g. memory or time limitations – to the derivations obtained from the descriptively adequate linguistic theory X, which is nothing but a formal description of our linguistic ‘Competence’. The distinction between Competence and Performance is traditionally viewed as logical and well-defined.¹¹ However, it is important to note that performance is the only empirical data source available to us for defining the boundaries of our language L. Thus, it is useful to introduce a four-way classification of performance data, independently considering the (i) meta-linguistic, (ii) categorical, (iii) temporal, and (iv) naturalistic dimensions. These dimensions, carved into the experimental perspective tradition, will enable us to frame the Language Problem (Definition 1) more effectively.

2.2.1. Meta-linguistic Dimension

Linguistic judgments can be ‘explicit’ – meta-linguistic – or ‘implicit’: we might ask a native speaker to judge, overtly, a sentence for grammaticality – or complexity, or acceptability – or ask the same speaker to simply read or repeat a sentence. In the first case, the speaker must access its meta-linguistic knowledge and provide a judgment that is otherwise unconscious unless overtly probed. In the second case, we can

consider implicit measures like reading-times or errors in the very same contexts.

2.2.2. Categorical Dimension

An observed linguistic behavior can be ‘categorical’ – binary – or ‘gradual’: if we ask for grammaticality judgments, a categorical judgment would be a binary choice between grammatical or ungrammatical, with no third option. The same judgment might be asked on a Likert scale from 1 (fully ungrammatical) to 7 (perfectly grammatical). An elicited answer to a specific linguistic question might also be considered categorical – correct or wrong – or gradual – receiving a score from 1 to 5, or similar. Although for a reasonable amount of repeated observations and a sufficient number of participants, these two measures converge (Sprouse & Almeida 2017), it remains a deep theoretical question to accept gradual judgments as a source of relevant empirical evidence for a theory. Considering, for instance, the generative power of a grammar (Chomsky 1956), categorical judgments are the only relevant empirical data we can consider.

2.2.3. Temporal Dimension

We can observe/elicit a linguistic behavior at the end of sentence processing (‘off-line’) or during sentence processing (‘on-line’): in the first case, we might ask a comprehension question or a grammaticality judgment when the exposure to the linguistic input to judge is completed. In the second case, we can record implicit – reading times in self-paced reading, fixation times in eye-tracking, neurometabolic responses recorded using EEG or fMRI, etc. – or explicit – e.g. maze techniques, Forster *et al.* 2009 – measures while the linguistic input under evaluation is processed, morpheme by morpheme, word by word or region by region.

2.2.4. Naturalistic Dimension

Naturalistic data are those collected using natural linguistic stimuli – e.g. narrative texts or public speeches – in ‘ecological’ conditions – e.g. reading or listening. This data collection approach, grounded in a solid corpus linguistic tradition, is complementary to the ‘experimental’ one, in which the linguistic input is accurately manipulated using minimal pairs and the task fully controlled to selectively evaluate the contribution of a specific variable in linguistic processing. Naturalistic/ecological studies are becoming very popular and are now widely considered useful sources of implicit processing information (Brennan *et al.* 2016; Siegelman *et al.* 2022). On the other hand, experimental practice on controlled inputs has significantly advanced our understanding of language

competence by targeting the influence of specific variables on minimal contrasts.

2.2.5. Between Experimental and Computational Linguistics

Employing a precise typology of empirical data assists us in reframing the tension between descriptive adequacy and computational efficiency, as well as clarifying the division of labor between competence and performance.

A standard practice, dating back to at least Chomsky and Miller's explanation of why native (naïve) speakers accept certain linguistic forms, (6a), but reject other structurally similar ones, (6b), is fundamentally tied to the competence/performance divide (Miller & Chomsky 1963):

(6) a. (I saw) [a dog [that bit [a cat [that chased [a mouse that ran away]]]]].
b. (I saw) [a mouse [that [a cat [that [a dog bit]] chased]] ran away].

A theory X, which predicts that restrictive relative clauses can modify indefinite Determiner Phrases (DPs), faces difficulty in restricting this (recursive) operation to non-nested contexts, such as (6a): opting for a solution that relies on an independent performance (algorithmic) domain, where available resources are limited and exhausted in (6b) but not in (6a), appears to be an appealing approach. As a result, our competence/computational perspective will be essentially declarative and make predictions on empirical data which are only (i) explicit (grammaticality judgments), (ii) categorical (grammatical vs ungrammatical), and (iii) off-line. Moreover, such an approach can hardly take advantage of naturalistic/ecological data while single grammatical rules/principles can be more easily tested under (iv) controlled conditions. While this option is tenable and logical, we can demonstrate that this solution is less descriptively adequate (Definition 2) when compared to a theory that can derive (D_s) a relevant structure (T_s) for a given sentence (S) with fewer assumptions, simply by operating at the algorithmic level.

Notice that, if we do not fully rely on native speaker judgments and we 'idealize' them, we widen the gap between the language as it's truly observed and the idealized one. By doing so, we implicitly adopt the following principle:

Definition 4. The dust under the carpet principle

If a theory X 'overgeneralizes' or 'undergeneralizes' on sentence S , we can maintain X by excluding the data related to performance on S .

i. a theory 'overgeneralizes' when it predicts a relevant structure T_s for S through derivation D_s , but native speakers judge S as ungrammatical because they cannot perform D_s ;

- ii. a theory ‘undergeneralizes’ when it is unable to predict a structure T_S for S through derivation D_S , but native speakers judge S as grammatical by performing D_S .

Everything being equal, it must be recognized that our descriptive adequacy definition suggests that a theory Y that can ‘remove the dust from under the carpet’ will be descriptively more adequate than X . Similarly, a theory Z that expands the empirical domain, taking advantage of many data sources, everything being equal, will be superior to X : by Definition 2, the theory that is able to generalize – i.e. capturing more data with fewer instructions – on more (i) explicit and implicit, (ii) categorical and gradual, (iii) off-line and online (iv) ecological and controlled data sources is descriptively more adequate.

In the end, the adopted notion of descriptive adequacy has the considerable advantage of being measurable: obviously, as clearly stated by Chomsky (Chomsky 2021a), we cannot simply rely on numerical comparisons since also the most trivial recursive grammar will be able to predict an infinite set of sentences/data. On the other hand, there is no need to weaken the idea that an ‘appropriate generalization’ is simply ‘a shorter set of instructions’ (or rules or principles) that capture an infinite set of (equivalent) constructions.¹² A theory that is equally descriptively adequate in capturing the same set of sentences at a comparable cost – number of ‘instructions’, equivalent structural T descriptions – will be challenged by different measures we can possibly obtain on the very same sentences: if a theory Z is not only able to predict that both sentences S_1 and S_2 are grammatical but also to which extent S_1 is processed more easily than S_2 , then theory Z must be considered descriptively superior compared to theory X which is not able to rank S_1 and S_2 .

So far, we have observed that: (i) although linguistic data available can be richly characterized, not all theoretical approaches can effectively utilize them; (ii) the trade-off between computational efficiency (fewer steps) and descriptive adequacy (fewer instructions) may favor larger theories that capture a broader array of data types. Consequently, the theories that are more descriptively adequate tend to be those that are observationally more adequate – that is, those that simply predict a larger quantity of data. This plainly supports Piantadosi’s point, which highlights the – observational – adequacy of vLLMs. However, in the end, this would lead to the inclusion of a vast number of irrelevant exceptions merely to broaden the empirical coverage of the theory. This approach would significantly disadvantage theories that cleverly apply the ‘dust under the carpet’ principle to dismiss a long tail of idiosyn-

cratic phenomena (e.g. those based on world knowledge rather than on linguistic competence).

Before precisely addressing an effective method for identifying ‘idiiosyncratic data’ (§3.1.2), it is crucial to mitigate a potential overemphasis on diverse empirical sources. To achieve this, it becomes essential to introduce an additional layer of adequacy based on the concept of ‘learnability’ (Chomsky 1965):

Definition 5. Explanatory adequacy

A theory that uses only primary linguistic data available to children to select a descriptively adequate model for L is explanatorily adequate.

Under Definition 5, if a theory Z only relies on positive primary linguistic data – a list of naturalistic Ss – to identify the language L (that also includes Ss), while a theory X needs ‘unreasonable’ data (as we will see in §3.1) to make the same predictions of Z , then Z is explanatorily adequate while X is not.

We are now ready to address two critical points with which generative grammar, as it is currently widely conceived in mainstream Minimalism, fails to comply. Those are related to, first, the formalism adopted and, second, a shared test set to be used as a benchmark.

Formalization Issue

If a theory is not fully explicit – i.e. formalized – there is no way to make precise predictions. For computational linguists accustomed to running their models on a computer, it is a well-known fact that no external oracle can ever fix a bug or a gap in the theory.

To the best of my knowledge, Edward Stabler was the first scholar to make a serious attempt at formally and successfully articulating a MG in a testable manner (Stabler 1997). His attempt was inspired by Chomsky’s 1995 formulation, but he filled various gaps to make the model sound and complete (in a pre-theoretical sense). One such example of ‘filling the gap’ was related to the definition of successful Merge and its impact on word order. Considering Merge as the fundamental operation that takes ‘a pair of syntactic objects (SO_i , SO_j) and replaces them by a new combined syntactic object SO_{ij} ’ (Chomsky 1995: 226), considering its asymmetric nature (Chomsky 1995: 246, either SO_i or SO_j will project) and the ‘inclusiveness condition’ (nothing but what is in the lexicon appears on phrase structure, Chomsky 1995: 249), Stabler formulated a feature-driven operation that strongly restricts the ‘free’ Merge operation: since we must limit the exuberance of this operation, before relying on later filters that

would reduce both descriptive adequacy and computational efficiency (under Definition 2 and Definition 3 respectively), feature checking is a suitable option: α and β , two independent syntactic objects, will (re-)merge successfully if and only if α has a *probe/select/interpretable/licensor* feature and β has a *goal/base/uninterpretable/licensee* feature.¹³ In more explicit terms, considering $=x$ a categorial probe associated with the head α and x a categorial goal associated with the complement β , then:

$$(7) \quad \text{Merge}(\alpha_{=x}, x\beta) = \{\alpha \{ \alpha_{=x}, x\beta \}\}$$

In prose: the Merge operation must take two items, α and β , that are characterized as the head (α , since it bears the probe feature ' $=x$ ', e.g. 'scolds_{=D}') and the complement (β , since it has the probed categorial feature ' x ', e.g. '_DBill'). The result of this – destructive, since features are deleted once the operation is successful – Merge operation is a T representation expressed in set-theoretical terms – α dominates/includes the unordered set $\{\alpha, \beta\}$, that is, $\{\text{scolds} \{ \text{scolds}_{=D}, {}_D\text{Bill} \}\}$. To obtain a predictable linear order, Stabler suggests that Merge probes the first goal 'to the right', while extra probes in α will be selected 'to the left' – this creates argumental shells (Stabler 1997: 7), e.g. in (3), (4): $\{\text{scolds} \{ {}_D\text{Alice}, \text{scolds}_{=D} \{ \text{scolds}_{=D =D}, {}_D\text{Bill} \} \}\}$; using angled brackets to indicate the linearization of words, the predicted outcome is: $\langle \text{Alice}, \text{scolds}, \text{Bill} \rangle$. The formalism developed by Stabler sparked a dynamic yet niche debate, primarily confined to the realm of mathematical logic and largely overlooked by both computational and generative linguists. This oversight persisted despite a commendable effort to update the model, incorporating newer concepts like Agree and the Workspace idea, through a collaboration with Chris Collins (Collins & Stabler 2016). Unfortunately, this update received limited attention, with only a select group of scholars recognizing the importance of a fully formalized theory (Chomsky, Seely *et al.* 2023).

Observe that these complex formalization practices are not mere mathematical curiosities. Instead, they represent the only tangible method for testing and refining linguistic theories. The original formulation of Merge lacked a crucial component, making the Language Problem (Definition 1) logically intractable. Without Stabler's proposal, this fundamental problem would have remained unresolved. Chomsky might regard this as a hybrid formalism that incorporates an externalization constraint – linear order. However, the inclusion of this constraint is a valid solution for generating sentences that would otherwise remain unobservable. Upon evaluation for descriptive adequacy, this methodol-

ogy facilitates direct comparison with other formalisms, enabling their ranking based solely on descriptive adequacy.

Furthermore, it is only through the formalizations provided by Stabler and his colleagues that we have been able to ascertain the expressive power of MG (Michaelis 2001) and explore different strategies for parsing and incremental analysis (Stabler 2013). Unfortunately, full theory formalization remains a rare practice in Minimalism – with laudable exceptions (Ginsburg & Fong 2019).

Evaluation Issue

Another essential component for fruitfully addressing the Language Problem is the creation of a shared test/reference set that encompasses all the relevant contrasts we aim to capture. Again, this is standard practice in computational linguistics: a model ‘is better’ if it performs better on a shared test set – according to various metrics. The relevant reference sets for MGs are already available as previously mentioned (SyntaxGym, §1). Another resource worth mentioning is the CoLA dataset – The Corpus of Linguistic Acceptability, Warstadt *et al.* 2018 – which comprises approximately 10K sentences from various linguistics publications, annotated for (binary) grammaticality.¹⁴ Similarly, the BLiMP dataset – a Benchmark of Linguistic Minimal Pairs for English, (Warstadt, Parrish *et al.* 2020) –, includes 1K minimal pairs, where each pair is evaluated by native speakers for preference (A is preferred over B). Notice that the size of the dataset is not the most important feature: as previously mentioned in note 11, if a theory X can accurately predict the structure of a transitive sentence where a relative clause modifies one DP, as illustrated in (6a), then it should also be capable of predicting an infinite number of similar sentences, where ‘similar sentences’ means those sentences that follow the same instructions, rules, or principles, despite variations in lexical items that are irrelevant to the structure. Therefore, the quality of a dataset, will not be strictly related to its size in general, but to the number of truly different structural configurations included. Moreover, a dataset that includes subtle attested contrasts, for instance, in ‘subject islands’ violations¹⁵ (8), (Chomsky 2008), will be more useful than one only including idealized ungrammatical cases (Huang 1982).

(8) a. *[Of which car]_i did [the driver _{..i}] cause a scandal?
b. [Of which car]_i was [the driver _{..i}] awarded a prize?

It is not an accident that Chomsky and colleagues apply the ‘dust under the carpet’ principle to islands by saying that “important judg-

ments can be quite murky. Without a clear understanding of such data, competing analyses can be difficult to compare" (Chomsky, Seely *et al.* 2023: 65). In fact, under Definition 2 and Definition 5, we have a practical way to assess theories that relegate 'under the carpet' a considerable corpus of experiments on islands constraints sensitivity (Sprouse & Hornstein 2013): descriptive inadequacy.

2.3. Objectives and Goals: an Intermediate Summary

To summarize, we can draw two relevant conclusions: (i) without a fully explicit (formalized) (Minimalist) theory, progress will be limited (formalization issue); (ii) without focusing on a shared and complete empirical domain – namely, a common test set – we cannot effectively compare the descriptive and explanatory adequacy of different formalized theories (evaluation issue).

I believe that the partial fulfillment of these requirements has caused generative linguistics to lose its footing and become marginalized in the contemporary landscape, which is dominated by efficient computational models, sophisticated experimental approaches, and, crucially, powerful statistical inferential methods serving both domains.

Let us go back to the three perspectives, to draw an interim conclusion: First, the computational perspective appears to be leading in terms of observational adequacy, and possibly in terms of descriptive adequacy as well, unless we consider to what extent the 'dust under the carpet' principle might fruitfully be used to exclude a long-tail or irrelevant idiosyncrasies. Moreover, while the computational perspective can address the learnability issue, it does so in a way that must be proven to be cognitively plausible (i.e. explanatory adequate). The theoretical perspective reminds us that restricting the empirical domain is necessary, due to the limitations imposed by our world knowledge, memory, and attention capacities. Unfortunately, by adopting this strategy, the most refined formalizations of Minimalist Grammar capture only a limited number of idealized contrasts, in comparison to the extensive body of available psycholinguistic evidence. Finally, the experimental perspective underscores that factors such as memory and attention, along with other performance confounds, can be quantitatively measured. Consequently, there is no justification for unduly restricting our empirical data-cake. As a result, psycholinguistic theories are increasingly adopting machine-learning methods over explicit generalizations, even though the latter offer the significant advantage of intelligibility.

To convince ourselves that linguistics is not merely a rock-paper-scissors game, we must revisit some foundational issues and appreciate

the advantage derived from an integration of these three different perspectives. The next section is devoted to the PoS hypothesis – a classic argument supporting the generative approach – and the notion of simplicity in structure-building operations, which is ostensibly a decisive factor in terms of learnability and evolvability. Indeed, I will argue that none of these issues can be adequately addressed from a singular perspective. Instead, all three viewpoints must be considered simultaneously to effectively tackle the Language Problem (Definition 1) and gain descriptive and explanatory adequacy.

3. Foundational Issues under Review

For a meaningful comparison of vLLMs and MGs that takes into account issues of formalization and evaluation, it is essential to revisit two foundational assumptions. These include the Poverty of Stimulus (PoS) hypothesis (§3.1) and the notion of simplicity in defining the core structure-building operation Merge (§3.2). The aim of this section is to underscore the importance of evaluating the impact of one specific formulation of structure-building operations against another. Ultimately, I will contend that although criticisms of vLLMs are warranted, they are frequently presented in an inappropriate or misguided manner.

3.1. The Poverty of Stimulus (PoS) Hypothesis: Learning from Positive Primary Linguistic Data

A cornerstone of generative linguistics is the PoS hypothesis, also known as Plato’s problem (Chomsky 1986). This hypothesis postulates the existence of innate knowledge to explain how children, under normal conditions and when sufficiently exposed to any natural language, can develop an adult-like linguistic knowledge – competence – that is not inferable from the limited exposure to the primary linguistic data available to them.¹⁶

Let us consider both a stronger and a weaker version of the PoS hypothesis: under the stronger interpretation, no adult linguistic competence can be attained, regardless of the amount of primary linguistic data available as input. Under the weaker interpretation, an adult’s linguistic competence can only be approximated after an ‘unreasonable’ – excessively large – exposure to primary linguistic data.

The stronger version is supported by classic results demonstrating the unlearnability (in the limit) of any recursive grammar, including the simplest regular ones (Gold 1967). These findings suggest that only trivial finite grammars are learnable when the primary linguistic input is

limited to positive evidence only, thus excluding negative evidence, that is, information indicating that a sentence is ungrammatical. Negative evidence is systematically excluded as a significant source of information for children because it is not uniformly provided – it is not given in all contexts or to all children – and is often noisy or inconsistent (Guasti 2017).

This limitation to positive primary linguistic inputs serves as a safeguard against machine learning methods employing supervised learning approaches. Such methods necessitate training on both well-formed and ill-formed sentences and implement a rewarding policy based on the model's accuracy in classifying grammatical versus ungrammatical sentences received as input. Compared to the ecological conditions of the primary linguistic input, those approaches require a training process that is clearly cognitively implausible, then irrelevant to address the 'logical problem of language acquisition' – 'PoS' in other words. Machine learning approaches, however, adopted another method that easily circumvents the safeguard banning supervised learning: Elman (Elman 1990) proposed a training trick, dubbed 'self-supervision', which is still used by modern Transformers (Vaswani *et al.* 2017) to train the most powerful vLLMs. When these models are trained, the only task they need to perform is to predict, as accurately as possible, the next word/token in S. Given a context such as *Please, open the...*, the next expected token should be 'something like' *door* or *window* and not like *is* or *John*. 'Something like', then, expresses a precise categorial generalization useful to provide a plausible token to successfully complete the next word prediction task. Simple Recurrent Neural networks (SRNs, Elman 1990) perform nicely on this task simply using, during training, the incoming word to check the accuracy of their prediction – therefore the term 'self-supervision'. This solution is (a) easy to implement, (b) psycholinguistically and cognitively plausible (Cloze probability, Bloom & Fischler 1980; Taylor 1953, priming effects, Bock 1986), (c) effective in capturing certain kinds of non-local dependencies (Elman 1993; Quinlan 2004). Finally, (d) it provides gradual predictions both off-line – an equivalent measure of the overall probability of generating a certain sequence S – and on-line – word-by-word, token-by-token preferences. It is, therefore, not surprising that this approach quickly gained popularity and reached a broad audience.

What is crucial to understand is that SRNs, among other capabilities, were able to learn certain non-local syntactic dependencies, such as subject-verb agreement, even when irrelevant lexical material (unrelated to the targeted agreement configuration) intervenes between the subject and the predicate. This is illustrated in (9) – as in (1):

(9) The dog [that the boys chase] grips/*grip the bone.

If trained SRNs prefer a third-person singular agreeing predicate over a non-agreeing one, regardless of the distance from the relevant subject, this indicates their observational adequacy in modeling such dependency. This preference suggests that SRNs implicitly prioritize grammatical structures over ungrammatical ones. Furthermore, this approach eliminates the distinction between idealized and externalized language (L) since it relies solely on observable (external) data to train and evaluate the model.

Many scholars have interpreted this as a significant challenge to the PoS hypothesis, suggesting that a tombstone may be placed over it. With a relatively simple ‘training program’ and using plausible primary linguistic inputs, a substantial portion of our grammatical knowledge can be acquired. Specifically, SRN-like networks exhibit human-like preferences in critical linguistic distinctions without depending on the explicit tree structures (T_s) associated with sentences (S_s), which are obtained through abstract derivations (D_s). Although Elman himself clarified that these findings did not disprove the Poverty of Stimulus (PoS) hypothesis – arguing that the network architecture, training procedure, and learning algorithm could all be considered forms of ‘innate knowledge’ (Bates *et al.* 1996) – the approach is generally perceived as a simpler, more elegant, and more appealing alternative to generative grammars and, subsequently, to Minimalism. This perspective has been reinforced by significant advancements: Long-Short Term Memory networks (LSTM, Hochreiter & Schmidhuber 1997) and their successors draw on Elman’s architectural and training insights and have introduced innovative solutions to address long-standing computational challenges, such as the vanishing gradient problem.¹⁷ Recently, networks of the LSTM type have been shown to outperform the widely acclaimed Transformer models (such as those used to train GPT-X vLLMs) in handling complex tasks within very subtle syntactic islands domains (Wilcox *et al.* 2024). To my knowledge, the paper by Wilcox *et al.* was the first to address various degrees of adequacy of these vLLMs in a typologically comprehensive manner on *Linguistic Inquiry*, which is probably the journal representing the most orthodox perspective on Minimalism.¹⁸ Similar findings led some generative linguists to adopt a weaker version of PoS: while these models might possess the capability to discern significant linguistic contrasts, they require an excessive amount of data for learning, especially when compared to the relatively limited data received by the children (Katzir 2023). I think accepting a weak perspective on PoS is detrimental to the generative cause. However, this position offers a valuable

opportunity to discuss several critical issues that the generative linguistic community has, perhaps neglectfully, adopted. Those lines of defense against vLLMs' descriptive and explanatory adequacy are generally linked to: (i) the excessive amount of data required to train these models (§3.1.1), (ii) the immense size of these models, which undermines their descriptive adequacy (§3.1.2), and (iii) the fact that language models lack an understanding of the meanings of the sentences they classify as grammatical (§3.1.3). I will address these arguments one by one in the following sections.

3.1.1. Training Data Size Argument

Many generative linguists may not yet be aware that, given a sufficiently large set of positive primary linguistic data – a size that can be reasonably quantified (Hsu & Chater 2010) – specific approaches, such as the Simplicity-Based framework (Hsu *et al.* 2013), allow for a relevant approximation of the ‘simplest’ grammar, namely shortest in MDL terms (§2.1).¹⁹ The approach adopted effectively reduces both overgeneralization and undergeneralization errors, which are respectively the result of overly general – simpler, shorter – or overly specific – more complex, longer – grammatical rules, by prioritizing a reduction in description length for grammar formulation – akin to a computer program’s length. Given that this framework can be logically implemented across various equivalent formalisms – e.g. mildly context-sensitive ones like Tree-Adjoining Grammars (Frank 1990, 2002), or MGs, (Stabler 2011, 2013) –, with no significant computational differences in terms of program length,²⁰ it aligns with the goal of descriptive adequacy (Definition 2) as defined within the Minimalist Program. While finding the exact shortest grammar is a non-computable problem (M. Li & Vitányi 2008), a probabilistic approximation makes this goal attainable. Thus, if we adopt the weakest version of the PoS hypothesis, the size of the training set might remain the only argument against the viability of vLLMs from a generative linguistics perspective.

From this perspective, we have pretty accurate predictions of the average number of words children are exposed to as primary linguistic input during their early years: from 3 to 11 million words per year (Hart & Risley 1992). A clear trade-off must then be considered: the complexity of a rule directly impacts the amount of data needed to integrate such a rule into our grammatical knowledge.²¹ So, approximation becomes a tricky issue here: is it sufficient to obtain a reasonably compact grammar that captures 99.9% of the data in a test set? I believe the answer depends on the dataset. If the test set accurately reflects a comprehensive range of linguistic configurations, especially those rarely found in

natural speech but documented by Stage of Acquisition (SoA), achieving this would be a significant accomplishment. However, if the test set is limited to statistically frequent or naturally occurring constructions, the value of such a result is comparatively limited. The key challenge lies in testing our internal (I-language) competence using an external (E-language) dataset – a corpus specifically designed to challenge theoretical assumptions.

Just to remain on ‘murky judgments’, let us again consider the subextraction of a *wh*- item from a direct object, (10a), and a context where, to avoid *wh*- subextraction from a complex NP (strong island), we rely on alternative structural solutions, like (10b). Can our learning procedure, which is essential for achieving explanatory adequacy, infer from very limited data that the avoidance of the application of a ‘rule’ – in (10b) – that applies to (10a), in fact, indicates a prohibition against configurations like (10b')?

(10) a. Who did you see a painting of _?
b. Did you see a painting representing whom exactly?
b'. *Who did you see a painting that was representing _ exactly?

Obviously, (10b') cannot be found in any naturalistic positive primary linguistic input. The only method to evaluate the theory’s ability to make relevant generalizations is by incorporating examples like (10b') into the shared test set. This means relying on any subtle empirical evidence available, crucially including explicit and controlled judgments in our test set. Although such examples may comprise only 0.1% of the test data, their ability to induce model prediction failures is telling. This would indicate that the seemingly accurate generalizations regarding (10a-b) are misleading. This conclusion stems from the realization that these generalizations rely on a non-uniform – or ‘genuine’ – interpretation of *wh*- dependencies, failing to account for the ungrammaticality of (10b'), as well as structurally equivalent infinite lexical variations of it.

From this viewpoint, arguing against the PoS hypothesis based solely on the size of the training set might not only be unnecessary but also counterproductive. A more compelling argument lies in the irrelevance of the volume of positive data received during training if a language model fails to generalize critical linguistic facts, such as the ungrammaticality of specific island violations, in line with robust native speakers’ judgments. Therefore, the careful selection of a shared test set, one that includes ungrammatical sentences, becomes crucial. Such a dataset is essential to demonstrate the enduring relevance of the PoS argument

and to show that the grammar inferred by the model deviates from human-like performance.

Consider now the classic debate on the unlearnability of auxiliary-subject inversion in English yes-no questions, addressed by Crain and Nakayama's experiment (Crain & Nakayama 1987). This experiment can be adapted to test various vLLMs to see if their strategy for forming 'yes-no questions' aligns with the structural solutions observed in children aged 3-5 years. Crain and Nakayama employed a straightforward elicitation task: providing a context targeting a character X, they prompted children to ask X a question like, 'Ask X if the boy who *is* watching Mickey Mouse *is* happy.' Despite various errors, children consistently avoided incorrect inversion of the embedded auxiliary with the matrix subject – i.e. sentences like '**is* the boy who *_is* watching Mickey Mouse *is* happy?' are never produced. This is a robust finding that persisted across tense changes – e.g. '...*is/was... was/is* happy' – and the substitution of an auxiliary with a modal – e.g. '*is/can... is* happy'. Interestingly, ChatGPT (based on the GPT-3.5 version as of May 2023) managed to correctly perform the matrix subject-auxiliary inversion without difficulty. This experiment might suggest that traditional arguments must be refined and that more evidence is needed to delve deeper into the PoS debate (§3.1.2).

Generative linguists exploring the implications of training language models with realistically plausible data volumes should examine the findings from a recent shared task, BabyLM (Warstadt *et al.* 2023) before fully embracing arguments based on training size. In summary, BERT-based optimized models (ELC-BERT, Charpentier & Samuel 2023) achieved an accuracy of approximately 0.85 (with 1 representing perfect accuracy) in distinguishing minimal pairs from the BLiMP dataset (Warstadt, Parrish *et al.* 2020) when trained on a plausible corpus of 100 million tokens. Their accuracy decreases to 0.80 when the training size is reduced to 10 million tokens. To be precise, achieving the mentioned accuracy required 450 epochs, meaning each sentence in the 100M token dataset was presented to the model 450 times.²² For the 10M token training, this number escalated to 4000 epochs, highlighting a significant efficiency concern. A child is a real 'few-shots learner' (Brown *et al.* 2020) in this respect. Further, two additional outcomes underscore the linguistic limitations of these models: firstly, the ELC-BERT system (Charpentier & Samuel 2023), despite winning the BabyLM competition, only achieved an accuracy of around 0.47 on the Mixed Signals Generalization Set (MSGs, Warstadt, Zhang *et al.* 2020) – a benchmark consisting of ambiguous binary classifications designed to test a model's preference for hierarchical versus surface-level generalizations – and

0.59 on CoLA (Warstadt *et al.* 2018). Secondly, while accuracy improves with more training, the models' alignment with online metrics, such as reading time, deteriorates (Steuer *et al.* 2023) suggesting a diminishing return on linguistic relevance with increased training. Putting all this evidence together, I conclude that the stronger version of the PoS argument remains unchallenged by pre-trained vLLMs currently tested.

Bear in mind that this conclusion is crucially supported by on-line performance data experimentally collected on subtle linguistic contrasts.

3.1.2. Descriptive Adequacy Argument and Data Problems

Connectionists (Rumelhart *et al.* 1999) heralded the transition from the traditional competence/performance distinction to a unified processing approach with considerable pride. The advantages appeared to be substantial: connectionist models can predict a wide range of data types, including (i) implicit and explicit, (ii) categorical and gradual, (iii) off-line and on-line, (iv) ecological and controlled. This significantly raises the benchmark for observational adequacy: given their capacity to incorporate a broader spectrum of data, connectionist theories have the potential to achieve greater observational adequacy. Once more, the generative linguistic safeguard is provided by Definition 2: the quality of data – encompassing generalization aspects beyond simple quantity, as discussed in sections §2.2.5 and §3.1.1 – and the principle that fewer instructions lead to better descriptive adequacy. This principle highlights a critical vulnerability of these models: their substantial size. The dimension of any sub-symbolic model can be expressed in terms of the number of parameters: this has nothing to do with linguistic parameters (Chomsky 1981), but it is a useful approximation of the model's learning dimensions. That is, a 10-parameter model has 10 dimensions to represent the problem. In addition to these dimensions, one must consider the length of the computer program required to utilize these parameters. However, due to the Invariance Theorem – note 20 –, this factor is typically negligible in most scenarios. In the context of modern vLLMs, GPT-3 is known to have 175 billion parameters. While the exact size of GPT-4 is not publicly disclosed (rumors indicate ~1.8 trillion parameters), the scaling hypothesis (Kaplan *et al.* 2020) implies a significant increase in parameters is necessary for improved performance. This represents a degree of complexity far surpassing that of any MG formalization.²³ Therefore, even when considering the Simplicity-Based account (Hsu *et al.* 2013), MGs are posited to surpass these vLLMs in the Descriptive Adequacy competition. This is obviously a wrong conclusion primarily because the full test set MGs can account for is significantly smaller than that addressed by vLLMs. Standard MG implementations I am aware of (Ginsburg & Fong 2019) are limited to

categorical grammaticality judgments, making them suitable for only a narrow range of predictions on naturalistic or gradual linguistic data. Incorporating statistical considerations (Hale 2001, 2011) is necessary to extend MGs applicability towards performance data. The situation marginally improves with the adoption of more processing-friendly MG formalisms (Chesi 2021a) or parsing-based metrics (De Santo 2019; Graf *et al.* 2017; J. Li & Hale 2019; Stabler 2013). These adaptations allow MG-based models to handle a broader variety of linguistic data types (§2.2.1–§2.2.4). Despite their enhanced ability to make linguistically informed predictions on minimal contrasts – a capability only recently tested against vLLMs (Hu *et al.* 2020) – the overall alignment of MG-based models with naturalistic data remains lower if compared to advanced surprisal measures (Futrell *et al.* 2020; Futrell & Levy 2017; Hale 2016; Levy 2008) based on robust statistical models.²⁴

Summarizing, vLLMs demonstrate remarkable robustness in harnessing diverse data sources obtained through sophisticated experimental methodologies, readily incorporating various empirical linguistic dimensions without hesitation. In contrast, formal models efficiently handle binary judgments but falter in capturing continuous, nuanced predictions. These theoretical models necessarily resort to the ‘dust under the carpet’ principle, ultimately marginalizing themselves from significant benchmark competitions. This is particularly evident in tests featuring relevant contrasts that do not pose any problem for a theory fragment taken in isolation but pose significant challenges to the overall theoretical framework consistency. Therefore, in challenging bulk testing scenarios, vLLMs exhibit a clear empirical advantage.

Data Dust: What Empirical Evidence Should We Ignore?

We are left with the uninterpretability of the implicit representations of these models, but this is, again, a hardly sustainable criticism. If a model’s description is shorter and thus descriptively more adequate, it ought to be preferred, as it is expected to capture essential generalizations despite their not being immediately intelligible. This dilemma aligns with the issues raised in Piantadosi’s paper, which argues that if two equations yield identical predictions with equal computational costs (in terms of steps taken or memory/dimensions utilized), they should be deemed equivalent in terms of descriptive adequacy.

In my view, however, there exists a compelling argument for favoring MG-based approaches over vLLMs on similar grounds. If vLLMs are indeed capable of encapsulating all relevant syntactic and semantic generalizations after relatively modest data exposure (between 10 million to 100 million tokens), with the remainder training corpus associated to

‘commonsense knowledge’ (Zhang *et al.* 2021), then we should narrow our focus to the modest set of parameters encoding structural generalizations, fruitfully adopting the ‘dust under the carpet’ principle. By doing so, we can undertake a precise comparison with MGs. This approach allows us to evaluate the efficiency and efficacy of each model in capturing the core linguistic (syntactic) phenomena, setting a more focused and fair ground for comparison. The rationale behind this comparison is that core syntactic properties ought to account for most data compression, leaving the idiosyncrasies to consume the bulk of the parameters in a vLLM. We can then attempt to define which types of data may be excluded from our test set by adopting a ‘Data Dust’ principle:

Definition 6. Data Dust

Irrelevant data for linguistic theorizing are those that force an increase in theory size without yielding any new generalizations

In more explicit terms, if a theory X for language L predicts a novel contrast C_n (e.g. a minimal sentence pair) without incrementing its size, then X already encompasses the necessary generalization. Thus, C_n becomes a valid candidate for inclusion in a test set for language L . Conversely, if no linguistic theory performs a correct prediction on C_n without increasing in size, and there are no other contrasts C_m that this size increase captures, then C_n qualifies as ‘data dust’. C_n can then be temporarily relegated ‘under the carpet’ until a new theory Y emerges that can predict it without increasing its size.

Measuring exactly the tolerable compression rate necessary to classify a contrast as relevant or not is out of the scope of this paper – this might be crucial for distinguishing linguistic competence from world knowledge –, but it seems to me clear that this perspective challenges the sustainability of probabilistic learning approaches as proposed, for example, by Hsu *et al.* (2013). Specifically, it posits that learners must cope with complex NP island constraints even in the absence of direct evidence for *wh*-item sub-extraction from a relative clause, a point elaborated in §3.1.1. From this viewpoint, the key is not merely accumulating a vast dataset but compiling a targeted collection of examples (possibly including minor lexical variations) that test the model’s handling of configurations unlikely to be found in naturalistic data. This approach is exemplified by the generation of pattern-based items in the AcCompl-IT dataset (Brunato *et al.* 2020), which are evaluated by human judges and all produce consistent responses (based on a 7-point Likert scale). Such phenomena include cases like illegally filled gaps (11) or unlicensed negative polarity items (12a), areas where the

performance of GPT-based models remains imperfect – they randomly accept gaps filled illegally, example (11) ending with *-lo*, and they license negative polarity items also in the absence of an appropriate negative licensor, example (12a) ending with *mai*.²⁵

(11) {Che cosa | quale problema}_i lo studente dovrebbe descriver(e) {_i|*_i-lo_i}?
{that what | which problem}_i the student must describe {_i|*_iit_i}
{what | which problem}_i must the student describe {_i | *_iit_i}

(12) a. Maria si aspetta che qualcuno possa avere {già | *mai}
M. himself expects that someone could have {already | *ever}
finito questo esercizio
completed this exercise
b. Nessuno si aspetta che qualcuno possa avere {già | mai}
Nobody himself expects that anyone could have {already | ever}
finito questo esercizio
completed this exercise
{M. | Nobody} expects {someone | anyone} to have {already/ever} finished this exercise.

I remind the reader that vLLMs are designed primarily for answering questions and executing various NLP tasks, rather than explicitly modeling human grammatical competence. These models, including GPT-4, represent a substantial advance in encapsulating a broad ontology – a detailed and structured body of knowledge that enables them to tackle challenges like the Winograd schema (Levesque 2014) and excel in numerous knowledge-based assessments (OpenAI 2023). Within this broad spectrum of capabilities, grammatical knowledge is merely one aspect under evaluation.

The performance of these models in dealing with complex grammatical distinctions, despite being trained for a general next-word prediction task, might seem remarkable. However, this capability is not entirely genuine. Supervised fine-tuning is crucial for transforming these systems into effective ‘few-shot learners’ (Brown *et al.* 2020). The nature and quality of this ‘fine-tuning’ – essentially a form of supervised learning – remain opaque. Consequently, when considering these models as ‘grammatical theories’ in the vein proposed by Piantadosi, their relevance becomes questionable. If the fine-tuning process leverages comprehensive datasets such as CoLA, it could render the learning challenge both biased and unreasonable, thereby diminishing the models’ relevance as cognitive models of language. This concern underscores my skepticism toward assessing linguistic competence by requesting grammatical judgments from ChatGPT.

On the Separation of PF from LF: A Challenge for Descriptive Adequacy

To conclude, returning to the T-model as outlined in (5), and the proposed division of labor between core syntax, PF, and LF, it is now crucial to reevaluate the extent to which this minimalist perspective provides a logical advantage over the integrated approach to processing, semantic, and syntactic knowledge offered by vLLMs.

A critical empirical challenge for the T-model, which cannot easily be swept under the carpet, concerns the interpretation of quantifiers and their scope. It has been proposed that Quantifier Raising (QR, May 1985) should be considered the ‘dust under the carpet’, namely a matter for the LF component. However, under certain very fruitful analyses, this operation is necessary to predict ‘extrapositions’, an operation that impacts on linear order – namely on the overt realization of S – as exemplified by the contrast (13b-b') below:

(13) a. I saw [{a (very good) / the (best)} picture [of the museum]] yesterday.
b. I saw [a (very good) picture _i] yesterday [of the museum]_i.
b'. ??I saw [the (best) picture _i] yesterday [of the museum]_i.

In a nutshell, the prepositional phrase [of the museum] can be ‘extraposed’ from its host, [a (very good) picture], in (13b), but crucially not in (13b'), due to the quantificational status of the determiner ‘a’ as opposed to ‘the’ (Baltin 2017; Fox & Nissenbaum 1999). If QR, which is supposed to happen optionally at LF, has an effect on linear order, which is supposed to be relevant only at PF, either we remove the optionality of QR from the theory (Beghelli & Stowell 1997) – but then how can we account for the optionality of extraposition? This is necessary, as suggested by the grammaticality of both (13a) and (13b) – or, otherwise, there is no clear way to relate LF phenomena with PF effects. Consequently, a theory that explicitly addresses these constraints, and integrates them within a unified structural framework without an increase in size (Definition 6) would offer a better descriptive adequacy than the orthodox T-model. Unfortunately, to the best of my knowledge, neither MGs nor vLLMs have been tested against these specific contrasts.

3.1.3. The Misunderstanding Argument

Despite their impressive ability to answer specific questions accurately, it is often highlighted that vLLMs do not truly ‘understand’ complex, well-formed sentences and can sometimes process incorrect ones – or ‘impossible’, in the sense described by Moro (Moro 2023; Moro *et al.* 2023). This criticism must be handled with care for precise reasons that are worth exploring in some detail.

Roni Katzir, for instance, responding to Piantadosi's paper (Katzir 2023), tried to elicit from ChatGPT paraphrases that require the correct structural analysis, as in the case of (6a) here repeated for convenience:

(6) a. (I saw) [a mouse [that [a cat [that [a dog bit]] chased]] ran away].

Katzir showed that some versions of the GPT model struggled to identify who was chasing whom. In my opinion, this is the only crucial test to determine whether a model's understanding of a sentence aligns with the assumed structural representation T or not. It is crucial to remember, however, that a comparison with human performance on these intricate test sentences remains necessary (§2.2.5). Another sound approach is to rely solely on the accuracy, stability, and convergence of judgments (Dentella *et al.* 2023).

Alternative strategies to prove the misrepresentation of T that are based on (i) illogical answers, (ii) incorrect metalinguistic explanations, or (iii) answers to ill-formed inputs are less effective. As far as (i) is concerned, asking logical or ethical reasoning (Chomsky, Roberts *et al.* 2023) can be useful to assess the social danger related to vLLMs, but does not provide any compelling evidence about the linguistic knowledge expressed by these models. This follows under the reasonable assumption that linguistic competence and other high cognitive functions (including reasoning, theory of mind, etc.) are independent segregated modules, in Fodor's sense (Fodor 1983), as explicitly highlighted also in Piantadosi's paper. As for (ii), asking for metalinguistic explanations (non-intuitive judgments) and then criticizing imprecise usage of morphosyntactic terms is extremely unfair: apart from linguistically educated native speakers, for instance, no one have a clear understanding of why the nominative case is generally associated with the grammatical subject in languages like English. As for (iii), asking ill-formed questions and collecting reasonable or unreasonable answers are equally weak approaches: standard Minimalism does not make any explicit assumption on this performance side. In other words, native speakers might indeed be capable of answering questions by identifying and correcting a subject-verb agreement error. However, understanding how this correction process works falls outside the computational scope of any MG as currently conceived. Similarly, it is well-documented that humans can easily misinterpret ill-formed, nonsensical questions (Bever & Townsend 2001), yet no MG tackles this issue. Following the set-theoretic tradition, Definition 1 (Language Problem), incorporates both positive and negative criteria – '*all and only* the sentences Ss belonging to language L'. The positive criterion ensures the model can capture any sentence in

language L, while the negative criterion excludes sentences not belonging to L. The negative restriction seems to be overlooked by vLLMs, which tend to be ‘collaborative’ with ill-formed inputs. Criticisms built on these arguments would just reinforce the idea that those models are more descriptively adequate (robust) than any MG since they also model how we could recover from ill-formed inputs by re-analysis. Proving that the recovery/re-analysis strategy is adequate or not is a matter for empirical investigation, requiring the collection of experimental data. It is important to stress, again, that mainstream MGs are unable to provide any such strategy.

3.1.4. Taking Stock of the PoS Hypothesis

To wrap up, in section §3.1, I argued that the PoS hypothesis remains unchallenged by results from current vLLMs, even though classic arguments might require refinement – this is the case, for instance, with auxiliary subject inversion in English polar questions, which appears to be too easy a task for vLLMs. I have also claimed that standard criticisms of these models, which are based on three principal arguments – training data size, model size, and the misunderstanding argument – are at best irrelevant and at worst ill-posed. A more substantial criticism relies on models’ ability to encompass infinite generalizations without increasing in size. Given that neither vLLMs nor MGs have been evaluated on this specific criterion, their descriptive and explanatory adequacy is yet to be determined. So far, we can conclude that vLLMs are observationally more adequate but lack explanatory adequacy. MGs are more compact and intelligible. This suggests that, when evaluated with a comparable test set, their descriptive adequacy could potentially be superior.

3.2. The Simplicity Mantra: Merge, Linear Order, and Cross-Linguistic Variation

To conclude this paper, I aim to discuss two remaining issues related to the core assumptions of Minimalism: the simplicity of structure-building operations and the external constraints influencing these operations – third factors (Chomsky, Seely *et al.* 2023). I think that these two fundamental issues contributed to the perception of the Minimalist program as an elusive framework.

3.2.1. Simplistic Structure-Building Operations

A little personal anecdote to introduce the issue: in 2002, I was preparing a seminar with Klaus Ables for Noam Chomsky and Danny Fox’s class at MIT on ‘foundational issues’. The topic assigned to us was

‘Vision and Language’. The reference list we received notably included Marr’s monumental work on Vision, unsurprisingly endorsing the generative grammar approach (Marr 1982: 28-29). During the preparation for our seminar, I had the opportunity for a brief interview with Tomaso Poggio (Marr & Poggio 1976; Riesenhuber & Poggio 1999). Then I inquired about the extent to which a Merge-like operation, described as “the simplest possible structure building operation, namely a set-formation operation”, could be integrated as a computational component in vision modeling. Poggio’s unequivocal reply is etched in my memory: “This does not make much sense to me, it’s too simplistic”. This could likely have been seen as a quick dismissal of an overeager PhD student’s question, but it persuaded me that, despite the fundamental necessity of an operation to create structure – which undoubtedly must be ‘as simple as possible’ – the Merge operation might be overly simplistic.

Without beating around the bush, an operation that overgenerates systematically is, computationally speaking, useless (Chomsky 1995: 283). It is easy to show that the unconstraint Merge operation used in (3) can produce whatever ungrammatical sentence we want – e.g. $\text{Merge}(\text{'scolds'}, \text{Merge}(\text{'Alice'}, \text{'Bill'})) = \{\text{scolds } \{\text{Alice, Bill}\}\}.$ ²⁶

Constraints come with a high cost, and due to Definition 2, a filter that excludes already generated unwanted structures will be discarded in favor of simpler (shorter, in description length terms) solutions. Although historically, a strategy involving filters was explored (Chomsky & Lasnik 1977), it is more convenient to pursue a more descriptively adequate solution that simply better constrains the derivation, thereby reducing unnecessary computation (Frampton & Gutmann 2002).

In their recent crystallization of Merge operation, Chomsky and colleagues (Chomsky, Seely *et al.* 2023) rely on Stabler & Collins formalism (Collins & Stabler 2016) and defend the set-formation idea behind this operation, as formulated in (3), under the simplicity and evolvability lens: the core operation to originate phrase structure must have been simple enough to evolve from a minimal genetic modification in our ancestral DNA. This line of reasoning does not extend to essential constraints like ‘labeling’, which involves selecting only the relevant information from the merged set that is useful for subsequent operations (Bošković 2016; Chomsky 2013; Rizzi 2016). Roughly speaking, labeling is necessary, for instance, to distinguish a verb phrase from a noun phrase without inspecting the content of the phrase built so far.

If labeling is not ‘part of Merge’ or something that must be learned, but rather a ‘third factor’ – a ‘natural law’ characterized by efficiency or ‘minimal search’ – then why haven’t we similarly recognized another widespread natural law like ‘incrementality’? Incrementality profoundly

influences our language performance, both in comprehension (Bever 1970) and production (Bock 1986). Given that phrase-structure building inherently progresses in (abstract) time, affecting every derivational approach crucially, it remains a mystery to me why this has not been recognized as a significant ‘law of nature’.²⁷

To appreciate this critical point let us return to the distinction between competence and performance. It logically follows that any performance-related task shall rely on one unique competence. Those tasks crucially include both comprehension and production – or parsing/recognition and generation, in computational terms, as overtly stated in the Language Problem (Definition 1). The fact that we need to draw upon the same competence theory is evident in the case of the lexicon: duplicating lexical entries to suit separate tasks for parsing and generation is highly inefficient. Yet, why has not this principle been equally applied to structure building? Merge operation, as conceived in (3), is incompatible with incremental parsing. This incompatibility arises because the order in which words are processed during sentence comprehension is the exact reverse – i.e., i. {Alice, scolds}, ii. {Alice, {scolds, Bill}} – of the order in which the structure (4) is derived according to (3) – i.e., i. {scolds, Bill} ii. {Alice, {scolds, Bill}}. As a consequence, a parsing/recognition algorithm must ‘undo’ Merge and Move, trying to guess the non-deterministic operations that might have resulted in specific word orders (Stabler 2013).

It is worth recalling that the original concept of perfection was aimed at finding an ideal solution for interface conditions (§2.1). Now, under the multiple spell-out approach, we have come to envision our core language faculty as a generative process that directs elements to PF and LF. Elements that, ironically, cannot be readily pronounced or (completely) interpreted!

This problem is easily illustrated for PF by example (6a), which is here repeated with an indication of the timing of the spell-out points associated with the CP phases:

(6) a. [_{ph 4} I saw a dog [_{ph 3} that bit a cat [_{ph 2} that chased a mouse [_{ph 1} that ran away]]]]].

Focusing solely on the complementizer layer and not taking other potential phasal domains into account, Phase 1 – ‘(mouse) that ran away’ – will be sent to PF before Phase 2 – ‘(cat) that chases a mouse [_{ph 1} (already spelled-out)]’. Such a derivation is logically possible, but it is completely non-sensical from a performance perspective. For the articulation of the complete sentence, PF should logically wait until the highest phase is complete before proceeding, but this approach is both counter-

intuitive and logically flawed. It is counterintuitive because speakers frequently commence utterances without having fully planned subsequent modifications, such as the addition of relative clauses towards the end of a sentence. It is logically flawed because it disregards the phenomenon of unbounded right recursion illustrated in (6a), which does not present a processing challenge. Such recursion demonstrates that the core syntactic engine inherently supports incrementality – producing sentences piece by piece. Unfortunately, this is not achievable with any mainstream MG that relies solely on Merge, as discussed in §2.1.

One might be misguided by the fact that incrementality seems to correspond to linear order, but this is an illusory perspective. Maintaining the premise that Merge merely constructs structure, as defined by (3), one might assume that:

- i. Merge is binary because of time ('third factor'): the incoming token (β) must be merged within the already formed structure (α);
- ii. Merge simply produces hierarchical structures, that is $\text{Merge}(\alpha, \beta) = [\alpha \, [\beta]]$

I fail to see how these two points contribute to an increased complexity of the Merge operation beyond that which is present in (3). However, I do recognize several benefits arising from this shift. The incorporation of incrementality considerations (Phillips 1996) represents a significant enhancement concerning interface conditions. In fact, this should be the null hypothesis (Momma & Phillips 2018) as it directly facilitates the delivery of syntactic objects that can be pronounced and interpreted incrementally. Consequently, it becomes possible to manipulate features that regulate nesting and movement,²⁸ thereby obviating the need for extraneous labeling considerations: the 'label' becomes the item that selects/expects another item as its complement. The alignment of this selection/expectation concept with the foundational principles of Elman's SRN and subsequent research is, in my view, remarkable. Determining whether specific artificial neural network architectures provide equivalent algorithmic solutions (in the sense proposed by Marr) for the computational and algorithmic predictions derived from this novel interpretation of Merge remains a question for empirical investigation. For example, the constraints associated with the application of this selection/expectation-based operation to (strong) islands (Bianchi & Chesi 2014) should emerge not from learned behavior but as a direct consequence of formalizing a structure-building operation that incorporates considerations on both nesting and incrementality. In this context, networks resembling SRN or LSTM appear more adept at reflecting the

concept of incrementality than those based on attention mechanisms. This observation aligns with empirical findings (Wilcox *et al.* 2024).

As discussed in §2.2.5, Stabler undertook the task of addressing numerous omissions in the original definition of Merge to delineate a well-defined structure-building operation capable of generating meaningful phrase structures. The introduction of feature checking, for instance, and the induction of linear order, may have been met with skepticism by some scholars. Nonetheless, these extensions were essential in rendering a particular variant of MG empirically testable. Other intuitions might well be formally articulated and empirically evaluated: If Merge is posited to be independent of feature checking, thus requiring a Labeling algorithm (Chomsky 2013, 2015), then it becomes imperative to demonstrate that derivations pertaining to the same test set can be accounted for in potentially more efficient and descriptive adequate ways than those offered by alternative feature-checking formulations. This evidence is currently missing especially in the case of (strong) islands.

3.2.2. Generative Parameters and Word Order Variation

After the Pisa lectures (Chomsky 1981), considerable research effort has been directed toward identifying a comprehensive list of parameters and organizing them into a coherent hierarchy (Baker 2001). This endeavor aims to render the problem of learnability more manageable. In this domain, I perceive the most significant advancements within generative linguistics: on the one hand, there has been a significant extension of the empirical basis, driven by radical generalizations regarding the restrictions of functional sequences (Cartographic approach). On the other hand, there has been a refinement in data elicitation methods that allows for an effective investigation into micro-parameterization, alongside the development of sophisticated mathematical methods for calculating phylogenetic distances (Gianollo *et al.* 2008; Guardiano *et al.* 2020; Guardiano & Longobardi 2016). A coherent picture is emerging in which parameters are not considered inherent components of Universal Grammar. Instead, they are viewed as options that remain underspecified until activated by specific selections of (lexical) feature bundles (Roberts 2019). Various resources sprouted from the cross-linguistic perspective, including atlases of different kinds. Notably, these include the World Atlas of Language Structures, WALS (Dryer & Haspelmath 2022), and the Syntactic Structures of the World's Languages, SSLW (Collins *et al.* 2009). Once again, the most straightforward method of engaging with heterogeneous resources involves

depending on the least ambiguous evidence they offer. For example, this can include preferences in word order expressed in terms of specific categories. An illustrative example is provided by Greenberg's Universal 20 (GU20), which focuses on word order within the extended nominal domain – DP. An original analysis emerged from Cinque's derivation of GU20 (Cinque 2005; Roberts 2017). Without digging into technical details, the attested and unattested word orders of four pertinent categories within the DP domain can be predicted by positing a universal hierarchical ordering of these categories, specifically [Dem [Num [Adj [NP]]]], plus a set of constraint on (head) movement. Cinque assumes Kayne's Linear Correspondence Axiom (LCA) to justify some of the constraints imposed on movement (Kayne 1994), while others not only consider this approach redundant but also dismiss LCA on the basis of an orthodox conception of Merge which expresses no order between the elements merged. According to this second view, if Merge operates on Adj and NP, then both possible linear orders, $\langle \text{Adj}, \text{NP} \rangle$ and $\langle \text{NP}, \text{Adj} \rangle$, could be observed across different languages (Abels & Neeleman 2012).

Both the approach relying on Kayne's LCA and the critique offered by Abels and Neeleman present plausible arguments that warrant thorough evaluation. While Abels and Neeleman's critique raises significant questions about the formal legitimacy of LCA, offering what appears to be a more elegant solution to the generalizations of GU20, an ultimate comparison of these proposals – considering simplicity, descriptive adequacy, and explanatory power – must be achieved. Moreover, an examination of the extensive empirical evidence assembled by Cinque, particularly with reference to data detailed in a dedicated section of the SSWL, uncovers instances of data idealization. This suggests that sweeping the 'dust under the carpet' remains a necessary principle for navigating through noisy data while still benefiting from coarse-grained categorial idealizations that can be later refined. While data idealization can accelerate early-stage research, the increasing precision of modern linguistic analysis necessitates a robust methodological approach. This approach should mirror established experimental practices and emphasize the creation of openly accessible resources. Such resources are essential for rigorous testing of any fully formalized theory.

In conclusion, formalizing fragments or assessing specific phenomena is essential for progress; however, in the absence of a unified framework in which to integrate coherently these fragments, the resulting overview appears as a Cubist patchwork. It is noteworthy that every complex discipline encounters similar crossroads: the Standard Model in quantum physics, for instance, is often criticized as a patchwork (Oerter

2006). Despite widespread dissatisfaction with this, it remains effective. It offers, to date, the most concise description of reality available, predicting phenomena with the highest accuracy possible. This includes phenomena that are rarely observable under natural conditions but can be artificially reproduced, such as those observed in the Large Hadron Collider (LHD, Evans 2007). By controlling the size, speed, and position of a few particles, researchers can capture better pictures of the collision events.

It appears to me that generative grammar remains anchored to classical – though solid – models, whereas computational and experimental linguistics have endeavored to make the quantum leap.

4. A Cautious Conclusion

At the end of the roundtable on Hilbert's List for Syntax, I had the impression that everyone was satisfied with the current problems' formulation, even though there was little interaction between one problem and another. I believe this reflects the sentiment that pervaded among those observing generative grammar from the outside: those idealized problems do not fit with each other and appear to be complex quirks with negligible impact on the understanding of the language faculty. The emerging trend associated with the computational and experimental perspectives is that concrete linguistic facts must be investigated using experimentally solid methods and computationally robust tools. Exotic theoretical puzzles with funny names and acronyms – e.g. strong islands, ATB extraction, Complex NP constraint – that are unattested in naturalistic corpora can be safely ignored if our aim is to understand and perform realistic linguistic tasks.

In these pages, I have attempted to demonstrate that this represents a limited perspective for achieving descriptive and explanatory adequacy. However, I am convinced that the current form of the generative paradigm, namely Minimalism, does not effectively support this position. This is due to a lack of consistency at both the formal level, where crucial intuitions that restrict structure-building operations remain underspecified, leading to formalization issues, and at the empirical level, where most theoretical intuitions are supported only by a limited empirical domain. Once this domain is extended, it often conflicts with other intuitions – evaluation issues.

In the end, on the one side of the field, computational linguists depend on statistical predictions obtained from vast corpora and have shown that the core syntactic engine, PF, and LF, are effectively distinct

only within the theoretical ‘T-model’. To truly understand what a sentence means – crucial for tasks like machine translation or answering questions – it is essential to rely on robust machine learning methods, which are more solid than any formal theory on the market. On the other side of the field, experimental linguistics has refined its methods, significantly improving the observational capabilities and, ultimately, enhancing the analysis of nearly all sources of linguistic data, whether implicit or explicit, categorical or gradual. Both players adopted open science practices, including the sharing of data and methods, and increasingly relied on sophisticated statistical methodology. Advanced statistical methods, as anticipated, have contributed to the success of machine learning and inferential analysis more than anything else. Generative linguists are sitting on the bench, watching the game, laughing at some experimental results – which seem occasionally to reinvent the wheel –, or expressing skepticism towards the inherent complexity of computational and statistical methods, as well as questioning their relevance as descriptive or explanatory theories. But they remained in the background. As Piantadosi provocatively said, this is “what happens when an academic field isolates itself from what should be complementary endeavours”. While generative linguistics struggles to accommodate gradual judgments, online effects, and other kinds of implicit data, these are the daily bread of computational and psycholinguistic models.

As demonstrated in these pages, generative linguistics has had the opportunity – and the ability – to establish the level of complexity for linguistic puzzles that need solving. However, this turned into a Pyrrhic victory: in focusing on the ‘speck’ in vLLMs’ eye, they failed to recognize the ‘log’ in their own, which, in my opinion, is the standard MG model’s inability to address incrementality, together with a lack of specification of fundamental constraints, and the absence of a shared evaluation benchmark.

While vLLMs are arguably overrated as linguistic theories, the methodology proposed by Wilcox and colleagues (Wilcox *et al.* 2024) represents an appropriate approach to testing them. By training different architectures from scratch on various plausible datasets (Warstadt *et al.* 2023), while avoiding fine-tuning, we may discover that the majority of the contrasts we aim to capture are indeed learnable. This outcome would offer evidence in favor of the PoS hypothesis, proving that the architectural intuitions explanatorily support language learnability. So far, the vLLMs tested are the only ones that perform properly on shared benchmarks such as SyntaxGym, CoLA, or BLiMP. Their dimension might be an issue, but only when a smaller model would obtain a comparable level of accuracy on these tests. In this respect, these vLLMs are,

in fact, really the best theories on the market, i.e. observationally more adequate than any MG.

It is widely recognized that new ideas – whether radical or minimal – are often supported by limited data, which computational and psycho-/neuro-linguists might deem marginal. The argument typically proposes a novel theoretical component, Y, needed to capture a distinction between evidence A and B – a distinction previously unaccounted for, thereby improving observational adequacy at the cost of increased size. Alternatively, it suggests replacing components Y and X with Z to accommodate the same dataset, thus enhancing descriptive adequacy through a reduction in theory size. Despite the innovation introduced by these ideas, their successful integration depends on addressing formalization and evaluation issues. The updated framework must undergo testing not only against the limited data that indicated the need for a novel component but also across the entire dataset. Such comprehensive evaluation is crucial to ensure that the new intuition neither overgeneralizes nor undergeneralizes.

I think the original sin of most generative linguists is that they have gotten used to incomplete pseudo-formalizations and data fragment explanations.

Making meaningful comparisons poses a significant challenge without addressing both the consistency and completeness of the theory, related to formalization issues, and establishing a reference dataset, pertaining to evaluation concerns. This situation has led to the rise of ‘Personal Minimalisms’, characterized by predominant subjective interpretations. Consequently, generative linguists may continue to superficially dismiss ‘murky judgments’. Meanwhile, psycho- and computational linguists could potentially uncover significant insights hidden within these ambiguities. We are well aware that not everything that glitters is gold. Ultimately, the most significant contribution that a generative linguist can provide is a linguistic minimal contrast challenging a specific theoretical assumption or the performance of a vLLM. Successfully incorporating this new contrast into a shared dataset, which any (r)evolutionary explicit formalism must confront, would represent quite a considerable accomplishment in my opinion.

If this does not happen, I fear that it might be the end of generative linguistics as we know it (but I feel fine).

Acknowledgements

This paper is partially supported by *NextGenerationEU* PRIN2022 grant, T-GRA2L (*Testing GRAdeness and GRAmmaticality in Linguistics*, 202223PL4N) and PRO₃ Scuole Superiori grant, BEXT (*New Behavioral EXperimental approaches to complexity perception and stress assessment in linguistic and cognitive research*) with the author as PI. The author wishes to thank Valentina Bianchi, Andrea Moro, Steven Piantadosi, plus the participants of the NeTS lab meetings for their insightful comments and valuable suggestions on an earlier draft of this paper. The author assumes full responsibility for any errors or omissions in this work.

Abbreviations

AoA = Age of Acquisition; ATB = Across-The-Board extraction (Williams 1977); BLiMP = Benchmark of Linguistic Minimal Pairs for English (Warstadt, Parrish *et al.* 2020); CI = Conceptual-Intentional interface; CoLA = Corpus of Linguistic Acceptability (Warstadt *et al.* 2018); CP = Complementized Phrase; DP = Determiner Phrase; GU20 = Greenberg's Universal 20 (Cinque 2005); LCA = Linear Correspondence Axiom (Kayne 1994); LF = Logical Form; LSTM = Long-Short Term Memory networks (Hochreiter & Schmidhuber 1997); MDL = Minimum Description Length (Rissanen 1978); MG = Minimalist Grammar; MSGS = Mixed Signals Generalization Set (Warstadt, Zhang *et al.* 2020); NP = Noun Phrase; PF = Phonological Form; PoS = Poverty of Stimulus; QR = Quantifier Raising (May 1985); SM = Sensory-Motor interface; SRN = Simple Recurrent Networks (Elman 1990); vLLM = very Large Language Models; VP = Verb Phrase.

Notes

¹ Adriana Belletti, Guglielmo Cinque, Denis Delfitto, Marcel Den Dikken, Anna Maria Di Sciullo, Robert Frank, Hubert Haider, Richard Kayne, Giuseppe Longobardi, M. Rita Manzini, Norvin Richards, Henk van Riemsdijk, Luigi Rizzi, Ian Roberts, Dominique Sportiche, and Peter Svenonius.

² The full list of abbreviations used is provided at the end of the paper.

³ In fact, a growing body of research suggests that the performance of state-of-the-art vLLMs falls short of human-level morphosyntactic competence. We will revisit this critical issue in §3.1 (Chesi *et al.* *to appear*; Dentella *et al.* 2023).

⁴ It is usually (anecdotally) assumed that sign languages pose a less strict requirement on lexical items linearization. The actual amount of simultaneity empirically observed in specific constructions, across sign languages is, in fact, precisely docu-

mented (Vermeerbergen *et al.* 2007).

⁵ I have deliberately used the term ‘instructions’ to include both the classic notion of (parametrized) principles – such as X'-theory, theta-criterion or case filter (Chomsky 1981) – and structure building operations – such as Merge or Move (Chomsky 1995).

⁶ One might be tempted to add “without generating unwanted structures (Frampton & Gutmann, 2002)” but this is already implicit in Definition 3: a theory Y that performs a wrong derivation D2, in addition to the appropriate one D1, is less efficient with respect to the theory X that only performs D1.

⁷ This is the size needed for a lexicon of just three words. Expanding the lexicon to 300 words, the code size will increase to approximately 5.7 KB. For a more extensive lexicon of 300,000 tokens, about 5.5 MB will be necessary. It is important to note that these lexical entries do not include any featural specifications and Select and Merge operations are completely unconstrained. For comparison, a vLLM such as GPT-3 requires about 350GB once optimized (Radford *et al.* 2018).

⁸ A ‘phase’ is a derivational chunk intended to minimize the computational burden. During the derivation, the items selected from the lexicon are arranged and sent to the interfaces – spelled-out – phase-by-phase. An uncontroversial phase – spell-out – domain is the Complementizer Phrase (CP), that is, embedded sentences are assembled and spelled-out before the matrix one.

⁹ A representational perspective only defines the relevant constraints in terms of the geometry of the syntactic tree. A classic representational constraint that applies to gaps/traces licensing is ‘C(onstituent)-command’: “Node A C-(onstituent)-commands node B if neither A nor B dominates the other and the first branching node which dominates A dominates B” (Reinhart 1976: 8). From a representational perspective the order in which single phrases are assembled is irrelevant. Conversely, this order is crucial in derivational approaches – MGs: from this perspective the order in which structure-building operations apply is fundamental to derive phrase structures. As a consequence, filters such as C-command must be reformulated in derivational terms, e.g. “X C-commands all and only the terms of the category Y with which X was paired/concatenated by Merge or by Move in the course of the derivation” (Epstein *et al.* 1998: 32).

¹⁰ This is in fact a pervasive solution to the motion problem in robotics (Nosengo 2014).

¹¹ A closer look might reveal logical inconsistencies: see the Competence Paradox (Chesi & Moro 2015).

¹² One might be tempted to provide a more explicit definition of ‘construction’, as Chomsky has done by referring to ‘legitimate generalizations’ (Chomsky 2021a). This is unnecessary: if a generalization T_S , obtained through D_S , is adequate – ‘legitimate’, in Chomsky’s terms –, regardless of how it is formulated, it will capture an infinite set of irrelevant lexical variations of the sentence S ; if it is inadequate – or ‘illegitimate’ –, it will fail to capture any variations. The issue then becomes one of size: the smaller, the better. This is what descriptive adequacy fundamentally entails (Definition 2).

¹³ I tried to include here all the relevant names which have been attributed to these features since Chomsky (1995). Stabler, in its original paper, used the ‘select’ vs ‘base’ terms for Merge, and ‘licensors’ and ‘licensee’ for Move. Here, I simplified a bit his formalism, for instance ignoring the directionality of the selection (to the right, $=X$ vs to the left, $X=$) and the strong/weak distinction ($=X$ and $=x$, respectively).

¹⁴ An equivalent resource is available for Italian: the ItaCoLA dataset (Trotta *et al.* 2021).

¹⁵ A ‘syntactic island’ is a phrasal domain from which extraction is not possible. The term was coined by John Ross (Ross 1967) and has been applied to various domains, including the Subject domain – example (8a) –, or the Complex NP (Noun Phrase) domain – ‘Complex NP constraint’ (Ross 1967: 127): “No element contained in a

sentence dominated by a noun phrase with a lexical head noun may be moved out of that noun phrase by a transformation". This is illustrated by the example (i) – the NP is currently labeled DP, the sentence CP, and the offending gap linked to the extracted NP 'the hat' (Ross 1967: 126):

(i) **The hat which I believed [_{DP} the claim [_{CP} that Otto was wearing _{the hat} is red]]*

¹⁶ It is important to emphasize that the term 'poverty' can refer to both 'quantity' – e.g. the number of words per year a child is exposed to – and 'quality' – e.g. a consistent lexical and structural diversity in linguistic input. In this discussion, I will primarily focus on the quantitative aspect, under the somewhat overoptimistic, yet reasonable, assumption that sufficient and consistent diversity is present in common child-directed speech.

¹⁷ In Recurrent Neural Networks that use standard backpropagation methods, we observe an exponential decrease of the error redistribution coefficient (i.e. the vanishing gradient) to be used to update the connection weights – i.e. learning – with, practically, no effect after few steps backward (Hochreiter *et al.* 2001).

¹⁸ Several other relevant tests had already been documented in the computational linguistics literature (Chowdhury & Zamparelli 2018; Linzen *et al.* 2016; Wilcox *et al.* 2018 a.o.) but were ignored by the generative linguistics community.

¹⁹ This, in fact, does not constitute the first challenge of Gold's unlearnability results (cf. Clark & Lappin 2010). The learnability of general rules (e.g. the 'pair' and 'append' operations on strings, which may encompass some key insights underlying the minimalist 'Merge') is also discussed in Y. Yang & Piantadosi (2022).

²⁰ This is the 'Invariance theorem' (M. Li & Vitányi 2008). In their experiment, Hsu & Chater (2010) employ a probabilistic context-free grammar, the adequacy of which we might reasonably question. The same experiment could be replicated using probabilistic MGs (J. Li & Hale 2019), potentially addressing phenomena more complex than the 19 contraction rules investigated in the original paper by Hsu & Chater.

²¹ Another way to formulate the trade-off between regular (simpler) rules and exceptions is the 'Tolerance Principle' (C. D. Yang 2016).

²² The situation is slightly more complex than it appears: a more accurate method for calculating data exposure should take into account the training steps, approximately 31k, and the batch size, about 8k, as reported in the appendix of the original paper for the 100M challenge submission. However, exploring the complexities of these concepts and the optimization steps falls outside the scope of this work.

²³ Including into a MG formalization both the lexicon and its features, using annotated Universal Dependencies treebanks (Nivre *et al.* 2017) to consider the most relevant kinds of lexical and syntactic ambiguity, we can estimate a lexicon size from about 13K distinct entries in English to about 33K in Turkish (Chesi 2023). Without considering efficient morphological decomposition (Kobele 2023), a full MG grammar is unlikely to require more than 100K parameter-equivalent dimensions. By approximation, we can equate this to the number of bits, as mentioned in note 7, which results in a size increase of at worst one order of magnitude – the compressed Python code, which includes the extracted English lexicon and is freely available at <github.com/cristianochesi/e-MGs>, has a size of merely 177KB. This size difference is substantial, spanning several orders of magnitude, especially when compared to the parameter scale of GPT models.

²⁴ A 'surprisal-like' measure is a value derived from information-theory metrics, indicating the unexpectedness of a word given its preceding context. The higher the surprisal value, the less likely the word is perceived as a natural continuation of the sentence. Such measures are utilized to predict processing difficulties in real-time reading or listening scenarios, with greater surprisal indicating higher expected difficulty. Additionally, surprisal values can be used to predict the outcomes of offline tasks: by calculating the cumulative surprisal for each word in a sentence, one can assess the overall unexpectedness of such sentence. For example, in forced-choice

tasks, the sentence with a lower total surprisal score is generally predicted to be more easily accepted or understood.

²⁵ In Italian, *mai* ('ever') is a negative polarity item which must be licensed by a C-commanding – see note 9 – negation like *nessuno* ('nobody'), as in the example (12a). Replacing *nessuno* with *Maria*, which is not a negative item, *mai* remains unlicensed. This explains the ungrammaticality of *Maria si aspetta che qualcuno possa avere *mai finito questo esercizio* ('M. expects someone to have ever finished this exercise').

²⁶ In this derivation '{Alice, Bill}' would be predicted to be a constituent, but this violates any reasonable syntactic test – e.g. *wh*-substitution of {Alice, Bill} 'constituent'. Moreover, the predicted linear orders for this sentence would be <scolds, A., B.>, <scolds, B., A.>, <A., B., scolds>, or <B., A., scolds>, none of which are attested, as base-generated, in standard English.

²⁷ Kayne explored the relevance of timing in his influential work (Kayne 1994), but to the best of my knowledge these considerations have not been pursued further.

²⁸ E.g. Complementation: Merge($\alpha_{\rightarrow x}, {}_x\beta$) = $[\alpha_{\rightarrow x} [{}_x\beta]]$ or Merge(${}_x\alpha, \beta_{=x}$) = $[[{}_x\alpha \beta_{=x}]]$ Free adjunction: Merge($\alpha_{\rightarrow x}, {}_y\beta$) = $[[{}_y\beta] \alpha_{\rightarrow x}]$ Movement: Merge($[[{}_y\beta]_i \dots \alpha_{\rightarrow x} [{}_x \gamma = {}_y \beta]]$, $[_y\beta]_i = $[[{}_y\beta]_i \dots \alpha_{\rightarrow x} [{}_x \gamma = {}_y \beta]]$$)

Bibliographical References

See the unified list at the end of this issue.

Unified Bibliographical References

Abels, Klaus & Neeleman, Ad 2012. Linear Asymmetries and the LCA: Linear Asymmetries and the LCA. *Syntax* 15,1. 25-74. <doi.org/10.1111/j.1467-9612.2011.00163.x>.

Abney, Steven 1996. Statistical methods. In Klavans, Judith L. & Resnik, Philip (eds.), *The Balancing Act: Combining Symbolic and Statistical Approaches to Language*. Cambridge, MA: MIT Press. 1-26.

Acemoglu, Daron 2024. *The Simple Macroeconomics of AI*. Working paper 32487. Cambridge, MA: National Bureau of Economic Research. <DOI: 10.3386/w32487>.

Achinstein, Peter 1985. *The Nature of Explanation*. Oxford: Oxford University Press.

Aksënova, Alëna & Deshmukh, Sanket 2018. Formal restrictions on multiple tiers. In *Proceedings of the society for computation in linguistics (SCiL) 2018*. 64-73.

Aksënova, Alëna; Graf, Thomas & Moradi, Sedigheh 2016. Morphotactics as tier-based strictly local dependencies. In *Proceedings of the 14th SIGMORPHON workshop on computational research in phonetics, phonology, and morphology*. 121-130.

Aksënova, Alëna; Rawski, Jonathan; Graf, Thomas & Heinz, Jeffrey 2024. The computational nature of hamony patterns. In Ritter, Nancy & van der Hulst, Harry (eds.), *Handbook of vowel harmony*. Oxford, UK: Oxford University Press. 437-451.

Allott, Nicholas; Kush, Dave & Dillon, Brian 2021. Sentence processing and syntactic theory. In Lohndal, T. & Rey, G. (eds.), *A Companion to Chomsky*. Wiley Publishing. 305-324.

Ambridge, Ben & Blything, Liam 2024. Large language models are better than theoretical linguists at theoretical linguistics. *Theoretical Linguistics* 50,1-2. 33-48.

Anderson, Chris 2008. The end of theory: The data deluge makes the scientific method obsolete. *Wired* 23 June.

Askell, Amanda; Bai, Yuntao; Chen, Anna; Drain, Dawn; Ganguli, Deep; Henighan, Tom; Jones, Andy; Joseph, Nicholas; Mann, Ben; DasSarma, Nova *et al.* 2021. A general language assistant as a laboratory for alignment. <arXiv:2112.00861>.

Ayers, John W. *et al.* 2023. Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. *JAMA Internal Medicine*. 589-596. <DOI: 10.1001/jamaintern-med.2023.1838>.

Baker, Mark 2001. *The atoms of language* (1st ed.). New York: Basic Books.

Baker, Mark 2009. Formal generative typology. In Heine, Bernd & Narrog, Heiko (eds.), *The Oxford Handbook of Linguistic Analysis*. 1st edition. Oxford: Oxford University Press. 285-312.

Baker, Mark 2013. On agreement and its relationship to case: Some generative ideas and results. *Lingua* 130. 14-32.

Baker, Mark 2021. On Chomsky's legacy in the study of linguistic diversity. In Allott, Nicholas; Lohndal, Terje & Rey, George (eds.), *A companion to Chomsky*. Hoboken, NJ: Wiley Blackwell. 158-171. <doi:10.1002/9781119598732.ch10>.

Baker, Mark & McCloskey, Jim 2007. On the relationship of typology to theoretical syntax. *Linguistic Typology* 11. 285-296.

Bai, Yuntao; Kadavath, Saurav; Kundu, Sandipan; Askell, Amanda; Kernion, Jackson; Jones, Andy; Chen, Anna; Goldie, Anna; Mirhoseini, Azalia; McKinnon, Cameron *et al.* 2022. Constitutional AI: Harmlessness from AI feedback. <arXiv:2212.08073>.

Baltin, Mark 2017. Extrapolosition. In Everaert, Martin & van Riemsdijk, Henk C. (eds.), *The Wiley Blackwell Companion to Syntax, Second Edition*. Hoboken, NJ: John Wiley & Sons, Inc. 1-33. <doi.org/10.1002/9781118358733.wbsyncom111>.

Barile, Joseph *et al.* 2024. Diagnostic accuracy of a Large Language Model in pediatric case studies. *JAMA Pediatrics*. 313-315. <DOI: 10.1001/jamapediatrics.2023.5750>.

Baroni, Marco 2022. On the proper role of linguistically oriented deep net analysis in linguistic theorizing. In Lappin, Shalom & Bernardy, Jean-Philippe (eds.), *Algebraic structures in natural language*. Boca Raton: CRC Press, Taylor & Francis. 1-16. *ICoRR* <arxiv.org/abs/2106.08694> (2021).

Barton, G. Edward; Berwick, Robert C. & Ristad, Eric Sven 1987. *Computational complexity and natural language*. Cambridge, MA: MIT Press.

Bates, Elizabeth; Elman, Jeffrey L.; Johnson, Mark H.; Karmiloff-Smith, Annette; Parisi, Domenico & Plunkett, Kim 1996. *Rethinking Innateness: A Connectionist Perspective on Development*. Cambridge, MA: MIT Press. <doi.org/10.7551/mitpress/5929.001.0001>.

Beghelli, Filippo & Stowell, Tim 1997. Distributivity and Negation: The Syntax of Each and Every. In Szabolcsi, Anna (ed.), *Ways of Scope Taking* (Vol. 65). Dordrecht: Springer Netherlands. 71-107. <doi.org/10.1007/978-94-011-5814-5_3>.

Beier, Eleonora J. & Ferreira, Fernanda 2022. Replication of Cutler, Anne & Fodor, Jerry A. 1979, Semantic focus and sentence comprehension. *Journal of Memory and Language* 126. <doi.org/10.1016/j.jml.2022.104339>.

beim Graben, Peter & Potthast, Roland 2014. Universal neural field computation. In Coombes, Stephen; beim Graben, Peter; Potthast, Roland & Wright, James (eds.), *Neural Fields*. Berlin: Springer. <doi.org/10.1007/978-3-642-54593-1_11>.

Belkin, Mikhail; Hsu, Daniel; Ma, Siyuan & Mandal, Soumik 2019. Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences* 116. 15849-15854. <doi.org/10.1073/pnas.1903070116>.

Bellelli, Adriana 2004. *Structures and Beyond: The Cartography of Syntactic Structures, Volume 3*. Oxford, UK: Oxford University Press.

Bender, Emily M.; Gebru, Timnit; McMillan-Major, Angelina & Shmitchell, Shmargaret 2021. On the dangers of stochastic parrots: Can language models be too big? New York, NY: Association for Computing Machinery. 610-623. <DOI: 10.1145/3442188.3445922>.

Bender, Emily M. & Hanna, Alex 2025. *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want*. Harper Collins.

Bender, Emily & Koller, Alexander 2020. Climbing toward NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the Association for Computational Linguistics*. 5185-5198. <www.aclweb.org/anthology/2020.acl-main.463>.

Benesty, Michaël 2023. *Unexpected description of GPT4 architecture*. <x.com/pommedeterre33/status/1671263789914677248>.

Bengio, Yoshua; Hinton, Geoffrey; Yao, Andrew; Song, Dawn; Abbeel, Pieter; Darrell, Trevor; Harari, Yuval Noah; Zhang, Ya-Qin; Xue, Lan; Shalev-Shwartz, Shai; Hadfield, Gillian; Clune, Jeff; Maharaj, Tegan; Hutter, Frank; Baydin, Atilim Gunes; McIlraith, Sheila; Gao, Qiqi; Acharya, Ashwin; Krueger, David; Dragan, Anca; Torr, Philip; Russell, Stuart; Kahneman, Daniel; Brauner, Jan & Mindermaann, Soren 2024. Managing extreme AI risks amid rapid progress. *Science* 384. 842-845. <doi.org/10.1126/science.adn0117>.

Berwick, Robert C. & Chomsky, Noam 2016. *Why only us: Language and evolution*. Cambridge, MA: MIT Press.

Berwick, Robert C.; Pietroski, Paul; Yankama, Beracah & Chomsky, Noam 2011. Poverty of the stimulus revisited. *Cognitive Science* 35,7. 1207-1242. <DOI: 10.1111/j.1551-6709.2011.01189.x>.

Bever, Thomas G. 1970. The cognitive basis for linguistic structures. *Cognition and the Development of Language*.

Bever, Thomas G. & Townsend, David J. 2001. Some Sentences on Our Consciousness of Sentences. In Dupoux, Emmanuel (ed.), *Language, Brain, and Cognitive Development: Essays in Honor of Jacques Mehler*. Cambridge, MA: MIT Press. 143-155.

Bianchi, Valentina & Chesi, Cristiano 2014. Subject islands, reconstruction, and the flow of the computation. *Linguistic Inquiry*. 525-569. <doi.org/10.1162/LING_a_00166>.

Bjorkman, Bronwyn M. 2017. Singular *they* and the syntactic representation of gender in English. *Glossa: A Journal of General Linguistics* 2,1. <DOI: 10.5334/gjgl.374>.

Blank, Idan 2016. *The Functional Architecture of Language Comprehension Mechanisms: Fundamental Principles Revealed with fMRI*. PhD dissertation. MIT. <doi.org/1721.1/7582>.

Bloom, Paul A. & Fischler, Ira 1980. Completion norms for 329 sentence contexts. *Memory & Cognition* 8,6. 631-642. <doi.org/10.3758/BF03213783>.

Bobaljik, Jonathan D. 2012. *Universals in comparative morphology: Suppletion, superlatives, and the structure of words*. Cambridge, MA: MIT Press.

Bobaljik, Jonathan D. & Wurmbrand, Susi 2008. Case in GB / Minimalism. In Malchukov, Andrej & Spencer, Andrew (eds.), *The Handbook of Case*. New York: Oxford University Press. 44-58.

Bobrow, Daniel G.; Cheslow, Bob; Condoravdi, Cleo; Karttunen, Lauri; Holloway King, Tracy; Nairn, Rowan; de Paiva, Valeria; Price, Charlotte & Zaenen, Annie 2007. PARC's bridge and question answering system. In *Proceedings of the Grammar Engineering Across Frameworks Workshop (GEFA 2007)*. CSLI Publications Online. 46-66.

Bock, J. Kathryn 1986. Meaning, sound, and syntax: Lexical priming in sentence production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12,4. 575-586. <doi.org/10.1037/0278-7393.12.4.575>.

Boeckx, Cedric & Leivada, Evelina 2013. Entangled parametric hierarchies: Problems for an overspecified Universal Grammar. *PLOS ONE* 8,9. <doi:10.1371/journal.pone.0072357>.

Bögel, Tina; Freiseis, Mila; Hill, Romi; Wambach, Daniel & Zhao, Tianyi 2024. Language redundancy and acoustic salience: An account in LFG. In Butt, Miriam; Findlay, Jamie A. & Toivonen, Ida (eds.), *The proceedings of the Ifg'24 conference*. 90-115.

Bögel, Tina & Zhao, Tianyi 2025. From speech signal to syntactic structure: A computational implementation. *Journal of Language Modeling* 13,1. 1-42.

Borer, Hagit 2005. *Structuring sense: In name only*. Oxford: Oxford University Press.

Bošković, Željko 2005. On the locality of left branch extraction and the structure of NP. *Studia Linguistica* 59. 1-45.

Bošković, Željko 2016. Introduction. *The Linguistic Review* 33,1. 1-16. <doi.org/10.1515/tlr-2015-0012>.

Bowman, Samuel R.; Hyun, Jeeyoon; Perez, Ethan; Chen, Edwin; Pettit, Craig; Heiner, Scott; Lukošiūtė, Kamilė; Askell, Amanda; Jones, Andy; Chen, Anna *et al.* 2022. Measuring progress on scalable oversight for large language models. <arXiv:2211.03540>.

Brayton, Flint; Laubach, Thomas & Reifschneider, David 2014. *The FRB/US Model: A Tool for Macroeconomic Policy Analysis*. Washington, DC: Board of Governors of the Federal Reserve System. <DOI: 10.17016/2380-7172.0012>.

Brennan, Jonathan R.; Stabler, Edward P.; Van Wagenen, Sarah E.; Luh, Wen-Ming & Hale, John T. 2016. Abstract linguistic structure correlates with temporal activity during naturalistic comprehension. *Brain and Language* 157-158. 81-94. <doi.org/10.1016/j.bandl.2016.04.008>.

Bresnan, Joan 1982. Control and complementation. *Linguistic Inquiry* 13,3. 343-434.

Bresnan, Joan 2016. Linguistics: The Garden and the Bush. *Computational Linguistics* 42,4. 599-617. <doi.org/10.1162/COLI a 00260>.

Bresnan, Joan; Cueni, Anna; Nikitina, Tatiana & Baayen, R. Harald 2007. Predicting the dative alternation. In Bouma, Gerlof; Krämer, Irene & Zwarts, Joost (eds.), *Cognitive Foundations of Interpretation*. Amsterdam: Royal Netherlands Academy of Science. 69-94.

Bressan, Veronica; Piccini Bianchessi, Maria Letizia; Fusco, Achille; Rossi, Sarah; Neri, Sofia & Chesi, Cristiano 2025. BLiMP-IT. <doi.org/10.17605/OSF.IO/2JKFN>.

Brown, Tom B.; Mann, Benjamin; Ryder, Nick; Subbiah, Melanie; Kaplan, Jared; Dhariwal, Prafulla; Neelakantan, Arvind; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; ... Amodei, D. 2020. Language Models are Few-Shot Learners. In Larochelle, Hugo *et al.* (eds.), *Advances in Neural Information Processing Systems 33 (NeurIPS 2020) Proceedings*. <arxiv.org/abs/2005.14165>.

Brunato, Dominique; Chesi, Cristiano; Dell'Orletta, Felice; Montemagni, Simonetta; Venturi, Giulia & Zamparelli, Roberto 2020. AcCompl-it@ EVALITA2020: Overview of the acceptability & complexity evaluation task for Italian. *Proceedings of Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2020), Online. CEUR. Org.*

Burness, Phillip; McMullin, Kevin & Chandlee, Jane 2021. Long-distance phonological processes as tier-based strictly local functions. *Glossa* 6. 1-37. <doi.org/10.16995/glossa.5780>.

Burness, Phillip; McMullin, Kevin & Nevins, Andrew 2024. Revisiting locality in vowel harmony. In Ritter, Nancy & van der Hulst, Harry (eds.), *Handbook of vowel harmony*. Oxford, UK: Oxford University Press. 269-290.

Butt, Miriam; Bögel, Tina; Zymla, Mark-Matthias & Mumtaz, Benazir 2024. Alternative questions in Urdu: from the speech signal to semantics. In Butt, Miriam; Findlay, Jamie & Toivonen, Ida (eds.), *Proceedings of the LFG'24 Conference*. Konstanz: PubliKon. 141-164. <lfg-proceedings.org/lfg/index.php/main/article/view/65/50>.

Butt, Miriam; Holloway King, Tracy; Niño, María-Eugenia & Segond, Frédérique 1999. *A Grammar Writer's Cookbook*. Stanford: CSLI Publications.

Butt, Miriam & Ramchand, Gillian 2005. Complex aspectual structure in Hindi/Urdu. In Ertishik-Shir, Nomi & Rappaport, Tova (eds.), *The Syntax of Aspect*. Oxford: Oxford University Press. 117-153.

Cahill, Aoife 2008. Treebank-based probabilistic phrase structure parsing. *Language and Linguistics Compass* 2,1. 36-58.

Cann, Ronnie; Kempson, Ruth & Marten, Lutz 2005. *The Dynamics of Language: An introduction*. Elsevier Academic Press.

Cao, Rosa & Yamins, Daniel 2024. Explanatory Models in Neuroscience, Part 2: Functional Intelligibility and the Contravariance Principle. *Cognitive Systems Research* 85. 101200. <doi.org/10.1016/j.cogsys.2023.101200>.

Carnie, Andrew 2013. *Syntax: A Generative Introduction, Third Edition*. Malden, MA: Wiley Blackwell.

Carnie, Andrew 2021. *Syntax: A Generative Introduction, Fourth Edition*. Malden, MA: Wiley Blackwell.

Cauchy, Augustin 1847. Méthode générale pour la résolution des systèmes d'équations simultanées. *Comptes rendus hebdomadaires des séances de l'Académie des sciences* 25. 536-538.

Cecchetti, Gabriele; Tomasini, Cedric A.; Herff, Steffen A. & Rohrmeier, Martin A. 2023. Interpreting rhythm as parsing. *Cognitive Science* 47. e13389. <doi.org/10.1111/cogs.13389>.

Chaitin, Gregory J. 1969. On the Simplicity and Speed of Programs for Computing Infinite Sets of Natural Numbers. *Journal of the ACM* 16,3. 407-422. <doi.org/10.1145/321526.321530>.

Chandlee, Jane 2014. Strictly local phonological processes. PhD dissertation. University of Delaware.

Chandlee, Jane 2017. Computational locality in morphological maps. *Morphology* 27. 599-641.

Chandlee, Jane 2022. Less is more: Reexamining assumptions through the narrow focus of subregularity. *Theoretical Linguistics* 48. 205-218.

Chandlee, Jane & Heinz, Jeffrey 2018. Strict locality and phonological maps. *Linguistic Inquiry* 49. 23-60.

Charchidi, Vincent J. 2024. Creative Minds Like Ours? Large Language Models and the Creative Aspect of Language Use. *Biolinguistics* 18. 1-31.

Charpentier, Lucas Georges Gabriel & Samuel, David 2023. Not all layers are equally as important: Every Layer Counts BERT. *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*. 210-224. <doi.org/10.18653/v1/2023.conll-babylm.20>.

Chen, Binglin; Lewis, Colleen M.; West, Matthew & Zilles, Craig 2024. Plagiarism in the age of Generative AI: Cheating method change and learning loss in an Intro to CS Course. In *L@S '24: Eleventh ACM Conference on Learning @ Scale, Atlanta GA USA*. New York, NY: ACM. 75-85. <DOI: [10.1145/3657604.3662046](https://doi.org/10.1145/3657604.3662046)>.

Chen, Tianlong; Frankle, Jonathan; Chang, Shiyu; Liu, Sijia; Zhang, Yang; Wang, Zhangyang & Carbin, Michael 2020. The lottery ticket hypothesis for pre-trained BERT networks. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F. & Lin, H. (eds.), *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*. Online: Curran Associates, Inc. 15834-15846.

Chen, Zhong & Hale, John T. 2010. Deforesting logical form. *Procs. Mathematics of Language*. Berlin: Springer. LNCS 6149. <doi.org/10.1007/978-3-642-14322-9_2>.

Cheng, Lisa L.-S.; Heycock, Caroline & Zamparelli, Roberto 2017. Two levels for definiteness. In Erlewine, M. Y. (ed.), *Proceedings of GLOW in Asia XI – Vol. 1. Volume 84 of MIT Working Papers in Linguistics*. MIT.

Cheng, Lisa L.-S. & Sybesma, Rint 1999. Bare and not-so-bare nouns and the

structure of NP. *Linguistic Inquiry* 30,4. 509-542.

Chesi, Cristiano 2007. An introduction to phase-based minimalist grammars: why move is top-down from left-to-right. In Moscati, V. (ed.), *STIL – Studies in Linguistics*, Volume 1. CISCL Press. 38-75.

Chesi, Cristiano 2021. Expectation-based Minimalist Grammars. <arxiv.org/abs/2109.13871>.

Chesi, Cristiano 2023. Parameters of cross-linguistic variation in expectation-based Minimalist Grammars (e-MGs). *Italian Journal of Computational Linguistics* 9,1. 21.

Chesi, Cristiano *forthcoming*. Linearization (as Part of Core Syntax). In Grohmann, Kleanthes & Leivada, Evelina (eds.), *Cambridge Handbook of Minimalism*. Cambridge (UK): Cambridge University Press. <ling.auf.net/lingbuzz/006689>.

Chesi, Cristiano; Barbini, Matilde; Bressan, Veronica; Neri, Sofia; Piccini Bianchessi, Maria Letizia; Sarah, Rossi & Sgrizzi, Tommaso 2024. Different Ways to Forget: Linguistic Gates in Recurrent Neural Networks. In *Proceedings of the BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*.

Chesi, Cristiano & Bianchi, Valentina 2014. Subject islands, reconstruction, and the flow of the computation. *Linguistic Inquiry* 45,4. 525-569.

Chesi, Cristiano & Moro, Andrea 2015. The subtle dependency between Competence and Performance. *MIT Working Papers In Linguistics* 77. 33-46.

Chesi, Cristiano; Vespignani, Francesco & Zamparelli, Roberto *to appear*. Large language models under evaluation: An acceptability, complexity and coherence assessment in Italian. *Italian Journal of Computational Linguistics*.

Chierchia, Gennaro 1998. Reference to kinds across languages. *Natural Language Semantics* 6. 339-405.

Cho, Kyunghyun; van Merriënboer, Bart; Gulcehre, Caglar; Bahdanau, Dzmitry; Bougares, Fethi; Schwenk, Holger & Bengio, Yoshua 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In Moschitti, Alessandro; Pang, Bo & Daelemans, Walter (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*. Doha, Qatar: Association for Computational Linguistics. 1724-1734. <DOI: 10.3115/v1/D14-1179>.

Chomsky, Noam 1956. Three models for the description of language. *IEEE Transactions on Information Theory* 2,3. 113-124. <doi.org/10.1109/TIT.1956.1056813>.

Chomsky, Noam 1957. *Syntactic Structures*. Berlin: Mouton de Gruyter.

Chomsky, Noam 1959. A Review of B. F. Skinner's Verbal Behavior. *Language* 35,1. 26. <doi.org/10.2307/411334>.

Chomsky, Noam 1964. *Current Issues in Linguistic Theory*. Berlin: De Gruyter.

Chomsky, Noam 1965. *Aspects of the Theory of Syntax* (Vol. 11). Cambridge, MA: MIT Press.

Chomsky, Noam 1966. *Cartesian Linguistics: A Chapter in the History of Rationalist Thought*. New York, NY: Harper & Row.

Chomsky, Noam 1968. *Language and Mind*. New York, NY: Harcourt, Brace & World.

Chomsky, Noam 1968b. Quine's Empirical Assumptions. *Synthese* 19,1-2. 53-68. <doi.org/10.1007/bf00568049>.

Chomsky, Noam 1969. Quine's empirical assumptions. In Davidson, Donald & Hintikka, Jaakko (eds.), *Words and Objections: Essays on the Work of W.V. Quine*. Dordrecht, Netherlands: Springer Dordrecht. 53-68. <DOI: 10.1007/978-94-010-1709-1_5>.

Chomsky, Noam 1975. *Questions on Form and Interpretation*. Lisse: Peter de Ridder. <doi.org/10.1007/978-3-642-14322-9_2>.

Chomsky, Noam 1981. *Lectures on government and binding: The Pisa lectures*. Walter de Gruyter.

Chomsky, Noam 1986. *Knowledge of language: Its nature, origin, and use*. New York: Praeger.

Chomsky, Noam 1995. *The minimalist program*. Cambridge, MA: MIT Press.

Chomsky, Noam 1995b. Language and Nature. *Mind* 104 (413). 1-61.

Chomsky, Noam 2001. Derivation by phase. In Kenstowicz, Michael (ed.), *Ken Hale: A life in language*. Cambridge, MA: MIT Press. 1-52.

Chomsky, Noam A. 2004. *The generative enterprise revisited. Discussions with Riny Huybregts, Henk van Riemsdijk, Naoki Fukui and Mihoko Zushi*. De Gruyter Mouton.

Chomsky, Noam A. 2005. Three Factors in Language Design. *Linguistic Inquiry* 36,1. 1-22.

Chomsky, Noam 2008. On phases. In Freidin, Robert; Otero, Carlos P. & Zubizarreta, Maria Luisa (eds.), *Foundational issues in linguistic theory: Essays in Honor of Jean-Roger Vergnaud* (Vol. 45). Cambridge, MA: MIT Press. 133-166.

Chomsky, Noam 2012. Language and Limits of Understanding. <www.nets.iusspavia.it/dox/chomsky2012-LLU-IUSS_Pavia.pdf>.

Chomsky, Noam 2013. Problems of projection. *Lingua* 130. 33-49.

Chomsky, Noam 2015. Problems of projection: Extensions. In Di Domenico, Elisa; Hamann, Cornelia & Matteini, Simona (eds.), *Linguistik Aktuell/Linguistics Today* (Vol. 223). Amsterdam: John Benjamins. 1-16. <doi.org/10.1075/la.223.01cho>.

Chomsky, Noam 2021a. Simplicity and the form of grammars. *Journal of Language Modelling* 9,1. <doi.org/10.15398/jlm.v9i1.257>.

Chomsky, Noam 2021b. Minimalism: where are we now, and where can we hope to go. *Gengo Kenkyu* 160. 1-42.

Chomsky, Noam 2024. The Miracle Creed and SMT. In Greco, M. & Moccia, D. (eds.), *A Cartesian dream: A geometrical account of syntax: In honor of Andrea Moro*. Rivista di Grammatica Generativa / Research in Generative Grammar 17-40.

Chomsky, Noam & Lasnik, Howard 1977. Filters and Control. *Linguistic*

Inquiry 8,3. 425-504.

Chomsky, Noam; Roberts, Ian & Watumull, Jeffrey 2023. Noam Chomsky: The False Promise of ChatGPT. *New York Times* 8 March.

Chomsky, Noam; Seely, T. Daniel; Berwick, Robert C.; Fong, Sandiway; Huybregts, M. A. C.; Kitahara, Hisatsugu; McInnerney, Andrew & Sugimoto, Yushi 2023. *Merge and the Strong Minimalist Thesis* (1st ed.). Cambridge: Cambridge University Press. <doi.org/10.1017/9781009343244>.

Chowdhury, Shammur Absar & Zamparelli, Roberto 2018. RNN Simulations of Grammaticality Judgments on Long-distance Dependencies. *Proceedings of the 27th International Conference on Computational Linguistics*. 133-144. <aclanthology.org/C18-1012>.

Cinque, Guglielmo 1999. *Adverbs and functional heads: A cross-linguistic perspective*. Oxford, UK: Oxford University Press.

Cinque, Guglielmo 2002. *Functional Structure in DP and IP: The Cartography of Syntactic Structures, Volume 1*. Oxford, UK: Oxford University Press.

Cinque, Guglielmo 2005. Deriving Greenberg's Universal 20 and Its Exceptions. *Linguistic Inquiry* 36,3. 315-332. <doi.org/10.1162/0024389054396917>.

Cinque, Guglielmo & Rizzi, Luigi 2010. The Cartography of Syntactic Structures. In Heine, B. & Narrog, H. (eds.), *The Oxford Handbook of Linguistic Analysis*. Oxford / New York: Oxford University Press. 51-65.

Clark, Alexander & Lappin, Shalom 2010. Computational learning theory and language acquisition. *Philosophy of Linguistics*. 445-475.

Clark, Alexander & Lappin, Shalom 2011. *Linguistic Nativism and the Poverty of the Stimulus*. Chichester: Wiley-Blackwell.

Clifton, Charles Jr; Ferreira, Fernanda; Henderson, John M.; Inhoff, Albrecht W.; Liversedge, Simon P.; Reichle, Erik D. & Schotte, Elizabeth R. 2015. Eye movements in reading and information processing. *Journal of Memory and Language* 86. 1-19.

Collins, Chris; Kayne, Richard & Koopman, Hilda 2009. *Syntactic structures of the world's languages (SSWL)*. <terraling.com/groups/7>.

Collins, Chris & Stabler, Edward P. 2016. A Formalization of Minimalist Syntax. *Syntax* 19,1. 43-78. <doi.org/10.1111/synt.12117>.

Collins, Joe 2024. The simple reason LLMs are not scientific models (and what the alternative is for linguistics). <lingbuzz.net/lingbuzz/008026>.

Conneau, Alexis; Kruszewski, German; Lample, Guillaume; Barrault, Loïc & Baroni, Marco 2018. What you can cram into a single \$&#!#* vector: Probing sentence embeddings for linguistic properties. In Gurevych, Iryna & Miyao, Yusuke (eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics. 2126-2136. <DOI: 10.18653/v1/P18-1198>.

Corbett, Greville G. 2010. Implicational hierarchies. In Song, Jae

Jong (ed.), *The Oxford Handbook of Linguistic Typology*. Oxford: Oxford University Press. 190-205. <doi.org/10.1093/oxfordhb/9780199281251.013.0011>.

Cottier, Ben; Rahman, Robi; Fattorini, Loredana; Maslej, Nestor; Besiroglu, Tamay & Owen, David 2025. The rising costs of training frontier AI models. <arXiv:2405.21015>.

Crain, Stephen & Nakayama, Mineharu 1987. Structure Dependence in Grammar Formation. *Language* 63,3. 522. <doi.org/10.2307/415004>.

Crain, Stephen & Thornton, Rosalind 2021. Universal grammar and language acquisition. In Allot, Nicholas; Lohndahl, Terje & Rey, Georges (eds.), *A Companion to Chomsky*. Wiley. <doi.org/10.1002/9781119598732.ch21>.

Crawford, Kate 2024. Generative AI's environmental costs are soaring – and mostly secret. *Nature* 626. 693. <DOI: 10.1038/d41586-024-00478-x>.

Crystal, David 2011. *Internet Linguistics: A Student Guide*. London: Routledge.

Cutler, Anne & Fodor, Jerry A. 1979. Semantic focus and sentence comprehension. *Cognition* 7. 49-59. <doi.org/10.1016/0010-0277(79)90010-6>.

Cybenko, George 1989 Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* 2,4. 303-314.

Dahl, Östen 2020. Morphological complexity and the minimum description length approach. In Arkadiev, Peter & Gardani, Francesco (eds.), *The complexities of morphology*. Oxford: Oxford University Press. 331-343.

D'Alessandro, Roberta 2019. The achievements of Generative Syntax: A time chart and some reflections. *Catalan Journal of Linguistics*. 7-26.

Dalrymple, Mary (ed.) 2023. *The Handbook of Lexical Functional Grammar: Empirically Oriented Theoretical Morphology and Syntax*. Berlin: Language Science Press. <10.5281/zenodo.10037797>.

Dalrymple, Mary; Gupta, Vineet; Lampert, John & Saraswat, Vijay 1999. Relating resource-based semantics to categorial semantics. In Dalrymple, Mary (ed.), *Semantics and syntax in Lexical Functional Grammar: The resource logic approach*. Language, Speech, and Communication. Cambridge, MA: MIT Press. 261-280.

Dalrymple, Mary; Patejuk, Agnieszka & Zymla, Mark-Matthias 2020. XLE + Glue – A new tool for integrating semantic analysis in XLE. In Butt, Miriam & Toivonen, Ida (eds.), *Proceedings of the LFG'20 Conference*. Stanford, CA: CSLI Publications. 89-108. <cslipublications.stanford.edu/LFG/2020/lfg2020-dpz.pdf>.

De Santo, Aniello 2019. Testing a Minimalist Grammar Parser on Italian Relative Clause Asymmetries. *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*. 93-104. <doi.org/10.18653/v1/W19-2911>.

De Santo, Aniello 2020. Structure and memory: A computational model of storage, gradience, and priming. PhD dissertation. Stony Brook University.

Deacon, Terence W. 1997. *The symbolic species: The co-evolution of language and the human brain*. Allen Lane: The Penguin Press.

Delétang, Grégoire; Ruoss, Anian; Grau-Moya, Jordi; Genewein, Tim; Wenliang, Li Kevin; Catt, Elliot; Cundy, Chris *et al.* 2022. Neural Networks and the Chomsky Hierarchy. <doi.org/10.48550/ARXIV.2207.02098>.

Demirci, Ozge; Hannane, Jonas & Zhu, Xinrong 2024. Who is AI replacing? The impact of Generative AI on online freelancing platforms. *SSRN Electronic Journal*. <DOI: 10.2139/ssrn.4991774>.

Demirdache, H.; Hornstein, N.; Lasnik, H.; May, R.; Rizzi, L. 2024. Structured Sentences and the Computational Theory of Mind: Roundtable. In *Festschrift for Howard Lasnik*. Cambridge: Cambridge University Press.

Dennett, Daniel C. 1978. Why you can't make a computer that feels pain. *Synthese* 38. 415-456.

Dentella, Vittoria; Günther, Fritz & Leivada, Evelina 2023. Systematic testing of three Language Models reveals low language accuracy, absence of response stability, and a yes-response bias. *Proceedings of the National Academy of Sciences* 120,51. e2309583120. <doi.org/10.1073/pnas.2309583120>.

Devlin, Jacob; Chang, Ming-Wei; Lee, Kenton & Toutanova, Kristina 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In Burstein, Jill; Doran, Christy & Solorio, Thamar (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Vol. 1. Minneapolis, MN: Association for Computational Linguistics. 4171-4186. <DOI: 10.18653/v1/N19-1423>.

Dijkstra, Edsger W. 1982. *Selected Writings on Computing*. Berlin: Springer.

Dobson, James E. 2023. On reading and interpreting black box deep neural networks. *International Journal of Digital Humanities* 5. 431-449. <DOI: 10.1007/s42803-023-00075-w>.

Dryer, Matthew S. 2006. Descriptive theories, explanatory theories, and basic linguistic theory. In Ameka, Felix K.; Dench, Alan & Evans, Nicholas (eds.), *Catching language: The standing challenge of grammar writing*. Berlin: Mouton de Gruyter. 207-234. <www.acsu.buffalo.edu/~dryer/desc.expl.theories.pdf>.

Dryer, Matthew & Haspelmath, Martin 2022. *The World Atlas of Language Structures Online* (v2020.3) [dataset]. Zenodo. <doi.org/10.5281/ZENODO.7385533>.

Edinger, Harald 2022. Offensive ideas: structural realism, classical realism and Putin's war on Ukraine. *International Affairs* 98,6. 1873-1893. <DOI: 10.1093/ia/iiac217>.

Elman, Jeffrey L. 1990. Finding Structure in Time. *Cognitive Science* 14,2. 179-211. <doi.org/10.1207/s15516709cog1402_1>.

Elman, Jeffrey L. 1991. Distributed representations, simple recurrent net-

works, and grammatical structure. *Machine Learning* 7,2. 195-225. <DOI: 10.1023/A:1022699029236>.

Elman, Jeffrey L. 1993. Learning and development in neural networks: The importance of starting small. *Cognition* 48,1. 71-99. <doi.org/10.1016/0010-0277(93)90058-4>.

Engelfriet, Joost; Lilin, Eric & Maletti, Andreas 2009. Extended multi bottom-up tree transducers: Composition and decomposition. *Acta Informatica* 46. 561-590. <doi.org/10.1007/s00236-009-0105-8>.

Epstein, Samuel David; Groat, Erich M.; Kawashima, Ruriko & Kitahara, Hisatsugu (eds.) 1998. *A derivational approach to syntactic relations*. Oxford, UK: Oxford University Press.

Ermolaeva, Marina 2023. Evaluating syntactic proposals using Minimalist grammars and minimum description length. *Journal of Language Modelling* 11. 67-119. <doi.org/10.15398/jlm.v11i1.334>.

Espinal, Maria Teresa & Cyrino, Sonia 2022. A syntactically-driven approach to indefiniteness, specificity and antispecificity in Romance. *Journal of Linguistics* 58. 535-570.

Ettinger, Allyson 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. <doi.org/10.1162/tacl_a_00298>.

Evans, Lyndon 2007. The Large Hadron Collider. *New Journal of Physics* 9,9. 335-335. <doi.org/10.1088/1367-2630/9/9/335>.

Evans, Nicholas & Levinson, Stephen C. 2009. The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences* 32,5. 429-448. <DOI:10.1017/S0140525X0999094X>.

Evanson, Linnea; Lakretz, Yair & King, Jean-Rémi 2023. Language acquisition: do children and language models follow similar learning stages? <arXiv:2306.03586>.

Fazi, M. Beatrice 2021. Beyond human: Deep learning, explainability and representation. *Theory, Culture & Society* 38. 55-77.

Feyerabend, Paul K. 1962. Explanation, reduction, and empiricism. In Feigl, Herbert & Maxwell, Grover (eds.), *Scientific explanation, space, and time*. Vol. 3. Minneapolis, MN: University of Minnesota Press. 28-97.

Fisher, Cynthia 2002. The role of abstract syntactic knowledge in language acquisition: A reply to Tomasello (2000). *Cognition* 82. 259-278.

Fleck, Ludwik 1935. *Entstehung und Entwicklung einer wissenschaftlichen Tatsache: Einführung in die Lehre vom Denkstil und Denkkollektiv*. Basel, Switzerland: Benno Schwabe & Co.

Fodor, Janet Dean 1998. Unambiguous triggers. *Linguistic Inquiry* 29. 1-36.

Fodor, Jerry A. 1980. *The Language of Thought*. Harvard: Harvard University Press.

Fodor, Jerry A. 1983. *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Fodor, Jerry A. 2010. *LOT 2: The Language of Thought Revisited*. Oxford, UK: Oxford University Press.

Fodor, Jerry A. & Bever, Thomas G. 1965. The psychological reality of linguistic segments. *Journal of Verbal Learning and Verbal Behavior* 4. 414-420. <doi.org/10.1016/s0022-5371(65)80081-0>.

Fong, Sandiway 1991. *Computational properties of principle-based grammatical theories*. PhD dissertation. MIT, Cambridge (MA).

Fong, Sandiway & Ginsburg, Jason 2012. Computation with doubling constituents: Pronouns and antecedents in Phase Theory. In Di Sciullo, Anna Maria (ed.), *Towards a Biolinguistic Understanding of Grammar: Essays on interfaces*. Amsterdam: John Benjamins. 303-338.

Fong, Sandiway & Ginsburg, Jason 2014. A new approach to tough-constructions. In Santana-LaBarge, Robert E (ed.), *Proceedings of the 31st West Coast Conference on Formal Linguistics (WCCFL 31)*. Somerville, MA: Cascadilla Proceedings Project. 180-188.

Fong, Sandiway & Ginsburg, Jason 2019. Towards a Minimalist Machine. In Berwick, Robert C. & Stabler, Edward P. (eds.), *Minimalist Parsing*. Oxford: Oxford University Press. 16-38.

Fong, Sandiway & Ginsburg, Jason 2023. On the computational modeling of English relative clauses. *Open Linguistics* 9. 1-35. <DOI: 10.1515/olip-2022-0246>.

Forster, Kenneth I.; Guerrera, Christine & Elliot, Lisa 2009. The maze task: Measuring forced incremental sentence processing time. *Behavior Research Methods* 41,1. 163-171. <doi.org/10.3758/BRM.41.1.163>.

Fox, Danny & Karzir, Roni 2024. Large Language Models and Theoretical Linguistics. *Theoretical Linguistics* 50. 71-76. <DOI: 10.1515/tl-2024-2005>.

Fox, Danny & Nissenbaum, Jon 1999. Extrapolation and scope: A case for overt QR. *Proceedings of the 18th West Coast Conference on Formal Linguistics* 18,2. 132-144.

Fox, Melvin J. & Skolnick, Betty P. 1975. *Language in Education: Problems and Prospects in Research and Teaching*. New York, NY: Ford Foundation.

Frampton, John & Gutmann, Sam 2002. Crash-Proof Syntax. In Epstein, Samuel David & Seely, T. Daniel (eds.), *Derivation and Explanation in the Minimalist Program* (1st ed.). Wiley. 90-105. <doi.org/10.1002/9780470755662.ch5>.

Frank, Anette; Holloway King, Tracy; Kuhn, Jonas & Maxwell, John T. III 2001. Optimality theory style constraint ranking in large-scale LFG grammars. In Sells, Peter (ed.), *Formal and Empirical Issues in Optimality Theory*. Stanford: CSLI Publications. 367-397.

Frank, Robert 1990. Licensing and tree adjoining grammar in government binding parsing. *28th Annual Meeting of the Association for Computational Linguistics*. 111-118.

Frank, Robert 2002. *Phrase structure composition and syntactic dependencies*. Cambridge, MA: MIT Press.

Frankle, Jonathan & Carbin, Michael 2019. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR 2019 Conference Track*. New Orleans, LA: OpenReview.

Friedmann, Naama; Belletti, Adriana & Rizzi, Luigi 2009. Relativized relatives: Types of intervention in the acquisition of A-bar dependencies. *Lingua* 119,1. 67-88.

Fusco, Achille; Barbini, Matilde; Piccini Bianchessi, Maria Letizia; Bressan, Veronica; Neri, Sofia; Rossi, Sarah; Sgrizzi, Tommaso & Chesi, Cristiano 2024. Recurrent Networks Are (Linguistically) Better? An Experiment on Small-LM Training on Child-Directed Speech in Italian. In *Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-It 2024)*. Aachen: CEUR.

Futrell, Richard; Gibson, Edward & Levy, Roger P. 2020. Lossy-Context Surprisal: An Information-Theoretic Model of Memory Effects in Sentence Processing. *Cognitive Science* 44,3. <doi.org/10.1111/cogs.12814>.

Futrell, Richard & Levy, Roger 2017. Noisy-context surprisal as a human sentence processing cost model. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. 688-698.

Futrell, Richard; Wilcox, Ethan; Morita, Takashi; Qian, Peng; Ballesteros, Miguel & Levy, Roger 2019. Neural language models as psycholinguistic subjects: Representations of syntactic state. <[arXiv:1903.03260](https://arxiv.org/abs/1903.03260)>.

Gauthier, Jon; Hu, Jennifer; Wilcox, Ethan; Qian, Peng & Levy, Roger 2020. SyntaxGym: An online platform for targeted evaluation of language models. In Celikyilmaz, Asli & Wen, Tsung-Hsien (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Online: Association for Computational Linguistics. 70-76. <DOI: 10.18653/v1/2020.acl-demos.10>.

Gehrke, Berit & McNally, Louise 2019. Idioms and the syntax/semantics interface of descriptive content vs. reference. *Linguistics* 57,4. 769-814. <[10.1515/ling-2019-0016](https://doi.org/10.1515/ling-2019-0016)>.

Gerth, Sabrina 2015. Memory limitations in sentence comprehension. A structure-based complexity metric of processing difficulty. PhD dissertation. University of Potsdam.

Gianollo, Chiara; Guardiano, Cristina & Longobardi, Giuseppe 2008. Three fundamental issues in parametric linguistics. In Biberauer, Theresa (ed.), *Linguistik Aktuell/Linguistics Today* (Vol. 132). Amsterdam: John Benjamins. 109-142. <doi.org/10.1075/la.132.05gia>.

Gibson, Edward; Futrell, Richard; Piantadosi, Steven T.; Dautriche, Isabelle; Mahowald, Kyle; Bergen, Leon & Levy, Roger 2019. How efficiency shapes human language. *Trends in Cognitive Sciences* 23,5. 389-407. <[10.1016/j.tics.2019.02.003](https://doi.org/10.1016/j.tics.2019.02.003)>.

Gibson, Edward & Wexler, Ken 1994. Triggers. *Linguistic Inquiry* 25,3. 407-454.

Gilkerson, Jill *et al.* 2017. Mapping the early language environment using

all-day recordings and automated analysis. *American Journal of Speech-Language Pathology* 26. 248-265. <DOI: 10.1044/2016_AJSLP-15-016>.

Ginsburg, Jason 2016. Modeling of Problems of Projection: A non-circular approach. *Glossa: A Journal of General Linguistics* 1,1:7. 1-46. <DOI: 10.5334/gjgl.22>.

Ginsburg, Jason 2024. Constraining free Merge. *Biolinguistics* 18, e14015. 1-60. <DOI: 10.5964/bioling.14015>.

Ginsburg, Jason & Fong, Sandiway 2019. Combining linguistic theories in a Minimalist Machine. In Stabler, Edward P. & Berwick, Robert C. (eds.), *Minimalist Parsing*. Oxford, UK: Oxford University Press. 39-68. <doi.org/10.1093/oso/9780198795087.003.0003>.

Giusti, Giuliana 2015. *Nominal Syntax at the Interfaces: A Comparative Analysis of Languages With Articles*. Cambridge: Cambridge Scholars Publishing.

Gold, E. Mark 1967. Language identification in the limit. *Information and Control* 10,5. 447-474. <doi.org/10.1016/S0019-9958(67)91165-5>.

Goldsmith, John & Riggle, Jason 2012. Information theoretic approaches to phonological structure: The case of Finnish vowel harmony. *Natural Language and Linguistic Theory* 30. 859-896.

Gorman, Kyle 2016. Pynini: A Python library for weighted finite-state grammar compilation. In *Procs. SIGFSM Workshop on Statistical NLP and Weighted Automata*. <doi.org/10.18653/v1/W16-2409>.

Goyal, Anirudh & Bengio, Yoshua 2022. Inductive Biases for Deep Learning of Higher-Level Cognition. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 478 (2266). <doi.org/10.1098/rspa.2021.0068>.

Graf, Thomas 2020. Curbing feature coding: Strictly local feature assignment. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2020*. 362-371.

Graf, Thomas 2022c. Typological implications of tier-based strictly local movement. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2022*. 184-193.

Graf, Thomas 2022b. Subregular linguistics: Bridging theoretical linguistics and formal grammar. *Theoretical Linguistics* 48. 145-184. <doi.org/10.1515/tl-2022-2037>.

Graf, Thomas 2022a. Diving deeper into subregular syntax. *Theoretical Linguistics* 48. 245-278. <doi.org/10.1515/tl-2022-2043>.

Graf, Thomas 2023. Subregular tree transductions, movement, copies, traces, and the ban on improper movement. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2023*. 289-299. <doi.org/10.7275/tk1n-q855>.

Graf, Thomas *to appear*. Minimalism and computational linguistics. In Grohman, Kleanthes K. & Leivada, Evelina (eds.), *Handbook of Minimalism*. Cambridge: Cambridge University Press.

Graf, Thomas & Abner, Natasha 2012. Is syntactic binding rational?

In *Proceedings of the 11th international workshop on Tree Adjoining Grammars and related formalisms (TAG + 11)*. 189-197.

Graf, Thomas & Kostyszyn, Kalina 2021. Multiple wh-movement is not special: The subregular complexity of persistent features in Minimalist grammars. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2021*. 275-285.

Graf, Thomas & Mayer, Connor 2018. Sanskrit n-retroflexion is input-output tier-based strictly local. In *Proceedings of SIGMORPHON 2018*. 151-160.

Graf, Thomas; Monette, James & Zhang, Chong 2017. Relative clauses as a benchmark for Minimalist parsing. *Journal of Language Modelling* 5.1. 57-106. <doi.org/10.15398/jlm.v5i1.157>.

Graf, Thomas & Shafiei, Nazila 2019. C-command dependencies as TSL string constraints. In Jarosz, Gaja; Nelson, Max; O'Connor, Brendan & Pater, Joe (eds.), *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*. 205-215.

Grice, Herbert Paul 1975. Logic and conversation. In Cole, Peter & Morgan, Jerry L. (eds.), *Syntax and Semantics*. New York, NY: Academic Press. 41-58.

Grillo, Nino 2008. *Generalized minimality: Syntactic underspecification in Broca's aphasia*. LOT.

Grünwald, Peter D. 2007. *The minimum description length principle*. Cambridge, MA: MIT Press.

Guardiano, Cristina & Longobardi, Giuseppe 2016. Parameter Theory and Parametric Comparison. In Roberts, Ian (ed.), *The Oxford Handbook of Universal Grammar*. Oxford, UK: Oxford University Press. 376-398. <doi.org/10.1093/oxfordhb/9780199573776.013.16>.

Guardiano, Cristina; Longobardi, Giuseppe; Cordoni, Guido & Crisma, Paola 2020. Formal Syntax as a Phylogenetic Method. In Janda, Richard D.; Joseph, Brian D. & Vance, Barbara S. (eds.), *The Handbook of Historical Linguistics* (1st ed.). Wiley. 145-182. <doi.org/10.1002/9781118732168.ch7>.

Guasti, Maria Teresa 2017. *Language acquisition: The growth of grammar*. Cambridge, MA: MIT Press.

Guérin, Jacqueline & May, Robert 1984. Extraposition and Logical Form. *Linguistic Inquiry* 15.1. 1-31.

Gulordava, Kristina; Bojanowski, Piotr; Grave, Edouard; Linzen, Tal & Baroni, Marco 2018. Colorless green recurrent networks dream hierarchically. In Walker, Marilyn; Ji, Heng & Stent, Amanda (eds.), *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. New Orleans, LA: Association for Computational Linguistics. 1195-1205. <[DOI: 10.18653/v1/N18-1108](https://doi.org/10.18653/v1/N18-1108)>.

Haider, Hubert 2023. Is Chat-GPT a grammatically competent informant? <lingbuzz/007285>.

Hale, John 2001. A Probabilistic Earley Parser as a Psycholinguistic Model.

Second Meeting of the North American Chapter of the Association for Computational Linguistics. <aclanthology.org/N01-1021>.

Hale, John 2011. What a rational parser would do. *Cognitive Science* 35,3. 399-443.

Hale, John 2016. Information-theoretical Complexity Metrics. *Language and Linguistics Compass* 10,9. 397-412. <doi.org/10.1111/lnc3.12196>.

Hanson, Kenneth 2025. Tier-based strict locality and the typology of agreement. *Journal of Language Modelling* 13,1. 43-97. <doi.org/10.15398/jlm.v13i1.411>.

Hanson, Kenneth 2024. Tiers, paths, and syntactic locality: The view from learning. In *Proceedings of the society for computation in linguistics (SCiL) 2024*. 107-116. <doi.org/10.7275/scil.2135>.

Hao, Sophie 2022. *Theory and Applications of Attribution for Interpretable Language Technology*. PhD dissertation. Yale University, New Haven, CT.

Hao, Sophie; Angluin, Dana & Frank, Robert 2022. Formal language recognition by hard attention transformers: Perspectives from circuit complexity. *Transactions of the Association for Computational Linguistics* 10. 800-810. <DOI: 10.1162/tacl_a_00490>.

Hao, Sophie; Mendelsohn, Simon; Sterneck, Rachel; Martinez, Randi & Frank, Robert 2020. Probabilistic predictions of people perusing: Evaluating metrics of language model performance for psycholinguistic modeling. In Chersoni, Emmanuele; Jacobs, Cassandra; Oseki, Yohei; Prévot, Laurent & Santus, Enrico (eds.), *Workshop on Cognitive Modeling and Computational Linguistics*. Online: Association for Computational Linguistics.

Hao, Sophie & Andersson, Samuel 2019. Unbounded stress in subregular phonology. In *Proceedings of the 16th Sigmorphon workshop on computational research in phonetics, phonology and morphology*. 135-143. <doi.org/10.18653/v1/W19-4216>.

Hart, Betty & Risley, Todd R. 1992. American parenting of language-learning children: Persisting differences in family-child interactions observed in natural home environments. *Developmental Psychology* 28,6. 1096-1105. <doi.org/10.1037/0012-1649.28.6.1096>.

Haspelmath, Martin 1993. *A grammar of Lezgian*. Mouton Grammar Library 9. Berlin: Mouton de Gruyter.

Haspelmath, Martin 2007. Pre-established categories don't exist – consequences for language description and typology. *Linguistic Typology* 11. 119-132.

Haspelmath, Martin 2008. Parametric versus functional explanations of syntactic universals. In Biberauer, Theresa (ed.), *The limits of syntactic variation*. Amsterdam: Benjamins. Accessed 27 May 2016.

Haspelmath, Martin 2010a. Comparative concepts and descriptive categories in crosslinguistic studies. *Language* 86,3. 663-687. <[doi:10.1353/lan.2010.0021](https://doi.org/10.1353/lan.2010.0021)>.

Haspelmath, Martin 2010b. Framework-free grammatical theory. In Heine,

Bernd & Narrog, Heiko (eds.), *The Oxford Handbook of Linguistic Analysis*. Oxford: Oxford University Press. 341-365.

Haspelmath, Martin 2018. How comparative concepts and descriptive linguistic categories are different. In Van Olmen, Daniël; Mortelmans, Tanja & Brisard, Frank (eds.), *Aspects of linguistic variation: Studies in honor of Johan van der Auwera*. Berlin: De Gruyter Mouton. 83-113. <zenodo.org/record/3519206>.

Haspelmath, Martin 2020. Human linguisticity and the building blocks of languages. *Frontiers in Psychology* 10,3056. 1-10. <[doi:10.3389/fpsyg.2019.03056](https://doi.org/10.3389/fpsyg.2019.03056)>.

Haspelmath, Martin 2021. General linguistics must be based on universals (or nonconventional aspects of language). *Theoretical Linguistics* 47,1-2. 1-31. <[doi:10.1515/tl-2021-2002](https://doi.org/10.1515/tl-2021-2002)>.

Haspelmath, Martin *to appear*. Breadth versus depth: Theoretical reasons for system-independent comparison of languages. In Nefdt, Ryan (ed.), *Oxford Handbook of Philosophy of Linguistics*. Oxford: Oxford University Press. <ling.auf.net/lingbuzz/008437>.

Hauser, M. D.; Chomsky, N. & Fitch, W. T. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298 (5598). 1569-1579. <doi.org/10.1126/science.298.5598.1569>.

Hawkins, John A. 2014. *Cross-linguistic variation and efficiency*. New York: Oxford University Press.

Heim, Johannes & Wiltschko, Martina 2025. Rethinking structural growth: Insights from the acquisition of interactional language. *Glossa: A journal of general linguistics* 10,1. <doi.org/10.16995/glossa.16396>.

Heinz, Jeffrey 2010. Learning long-distance phonotactics. *Linguistic Inquiry* 41. 623-661. <doi.org/10.1162/LING_a_00015>.

Heinz, Jeffrey 2018. The computational nature of phonological generalizations. In Hyman, Larry & Plank, Frank (eds.), *Phonological typology*. Mouton De Gruyter. 126-195.

Hewitt, John & Manning, Christopher D. 2019. A structural probe for finding syntax in word representation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 4129-4138.

Hey, Tony; Tansley, Stewart; Tolle, Kristin & Gray, Jim (eds.) 2009. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Redmond, WA: Microsoft Research.

Hinton, Geoffrey 2022. The forward-forward algorithm: Some preliminary investigations. <[arXiv:2212.13345](https://arxiv.org/abs/2212.13345)>.

Hochreiter, Sepp; Bengio, Yoshua; Frasconi, Paolo & Schmidhuber, Jürgen 2001. Gradient flow in recurrent nets: The difficulty of learning long-term dependencies. In Kremer, S. C. & Kolen, J. F. (eds.), *A Field Guide to Dynamical Recurrent Neural Networks*. IEEE Press.

Hochreiter, Sepp & Schmidhuber, Jürgen 1997. Long short-term memory. *Neural Computation* 9,8. 1735-1780.

Hockenmaier, Julia & Steedman, Mark 2007. CCGbank: A corpus of CCG derivations and dependency structures extracted from the Penn Treebank. *Computational Linguistics* 33,3. 355-396. <DOI: 10.1162/coli.2007.33.3.355>.

Holmes, V. M. & Forster, K. 1972. Click location and syntactic structure. *Perception and Psychophysics* 12. 9-15. <doi.org/10.3758/bf03212836>.

Hornik, Kurt; Stinchcombe, Maxwell & White, Halbert 1989. Multilayer Feedforward Networks Are Universal Approximators. *Neural Networks* 2,5. 359-66. <[doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)>.

Hosseini, Eghbal A. *et al.* 2024. Artificial neural network language models align neurally and behaviorally with humans even after a developmentally realistic amount of training. *Neurobiology of Language*. Apr 1.5,1. 43-63.

Hsu, Anne S. & Chater, Nick 2010. The Logical Problem of Language Acquisition: A Probabilistic Perspective. *Cognitive Science* 34,6. 972-1016. <doi.org/10.1111/j.1551-6709.2010.01117.x>.

Hsu, Anne S.; Chater, Nick & Vitányi, Paul 2013. Language Learning From Positive Evidence, Reconsidered: A Simplicity-Based Approach. *Topics in Cognitive Science* 5,1. 35-55. <doi.org/10.1111/tops.12005>.

Hu, Jennifer; Gauthier, Jon; Qian, Peng; Wilcox, Ethan & Levy, Roger 2020. A Systematic Assessment of Syntactic Generalization in Neural Language Models. In Jurafsky, Dan; Chai, Joyce; Schluter, Natalie & Tetreault, Joel (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics. 1725-1744. <doi.org/10.18653/v1/2020.acl-main.158>.

Hu, Michael Y.; Mueller, Aaron; Ross, Candace; Williams, Adina; Linzen, Tal; Zhuang, Chengxu; Cotterell, Ryan; Choshen, Leshem; Warstadt, Alex & Wilcox, Ethan 2024. Findings of the Second BabyLM Challenge: Sample-Efficient Pretraining on Developmentally Plausible Corpora. <doi.org/10.48550/ARXIV.2412.05149>.

Huang, C.-T. James 1982. *Logical relations in Chinese and the theory of grammar*. Cambridge, MA: MIT Press.

Huang, Lei; Yu, Weijiang; Ma, Weitao; Zhong, Weihong; Feng, Zhangyin; Wang, Haotian; Chen, Qianglong; Peng, Weihua; Feng, Xiaocheng; Qin, Bing *et al.* 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. <[arXiv:2311.05232](https://arxiv.org/abs/2311.05232)>.

Huh, Minyoung; Cheung, Brian; Wang, Tongzhou & Isola, Phillip 2024. Position: The Platonic Representation Hypothesis. In Salakhutdinov, Ruslan; Kolter, Zico; Heller, Katherine; Weller, Adrian; Oliver, Nuria; Scarlett, Jonathan & Berkenkamp, Felix (eds.), *Proceedings of the 41st International Conference on Machine Learning*. 235. 20617-42. Proceedings of Machine Learning Research. PMLR. <proceedings.mlr.press/v235/huh24a.html>.

Hume, David 1739. *A Treatise of Human Nature: Being an Attempt to Introduce*

the Experimental Method of Reasoning Into Moral Subjects. London: John Noon.

Hume, David 1748. *Philosophical Essays Concerning Human Understanding*. London: A. Millar.

Hunter, Tim; Stanojević, Miloš & Stabler, Edward P. 2019. The active-filler strategy in a move-eager left-corner Minimalist grammar parser. In *Proceedings of the workshop on cognitive modeling and computational linguistics*. 1-10.

Ibbotson, Paul & Tomasello, Michael 2016. Evidence rebuts Chomsky's theory of language learning. *Scientific American* 315, 5. 70.

İdrisoğlu, İşıl & Spaniel, William 2024. *Information problems and Russia's invasion of Ukraine*. *Conflict Management and Peace Science* 41,5. 514-533. <DOI: 10.1177/07388942241238583>.

Ionin, Tania & Matushansky, Ora 2006. The composition of complex cardinals. *Journal of Semantics* 16. 315-360.

Jackendoff, Ray 1988. Why are they saying these things about us? *Natural Language and Linguistic Theory* 6,3. 435-442.

Jardine, Adam 2016. Computationally, tone is different. *Phonology* 33. 247-283. <doi.org/10.1017/S0952675716000129>.

Ji, Zwei; Lee, Nayeon; Frieske, Rita; Yu, Tiezheng; Su, Dan; Xu, Yan; Ishii, Etsuko; Bang, Ye Jin; Madotto, Andrea & Fung, Pascale 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys* 55,12. 248:1-248:38. <DOI: 10.1145/3571730>.

Jurafsky, Dan & Martin, James H. 2008. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 2nd edition. Russell, Stuart & Norvig, Peter (eds.). Upper Saddle River, NJ: Prentice Hall.

Kalouli, Aikaterini-Lida 2021. *Hy-NLI: A hybrid system for state-of-the-art natural language inference*. University of Konstanz dissertation.

Kalouli, Aikaterini-Lida; Crouch, Richard & de Paiva, Valeria 2020. Hy-NLI: A hybrid system for natural language inference. In *Proceedings of the 28th International Conference on Computational Linguistics*. Barcelona, Spain (Online): International Committee on Computational Linguistics. 5235-5249. <aclanthology.org/2020.coling-main.459>.

Kaplan, Jared; McCandlish, Sam; Henighan, Tom; Brown, Tom B.; Chess, Benjamin; Child, Rewon; Gray, Scott; Radford, Alec; Wu, Jeffrey & Amodei, Dario 2020. *Scaling Laws for Neural Language Models*. <doi.org/10.48550/ARXIV.2001.08361>.

Kaplan, Ronald M. 1987. Three seductions of computational linguistics. In Whitelock, P.; Wood, M. M.; Somers, H.; Johnson, R. & Bennett, P. (eds.), *Linguistic Theory and Computer Applications*. London: Academic Press. 149-188.

Kaplan, Ronald M. 2019. Computational psycholinguistics. *Computational Linguistics* 45,4. 607-626. <doi:10.1162/coli_a_00359>. <aclanthology.org/J19-4001>.

Kaplan, Ronald M.; King, Tracey H. & Maxwell, John T. III 2002. Adapting

existing grammars: The XLE experience. In *COLING-02: Grammar Engineering and Evaluation*.

Katz, Phillip 1986. PKZIP. Commercial Compression System, Version 1.1. <www.pkware.com/pkzip>.

Katzir, Roni 2023. *Why large language models are poor theories of human linguistic cognition. A reply to Piantadosi (2023)* [LingBuzz]. <[lingBuzz/007190](https://lingBuzz.org/007190)>. *Biolinguistics* 17. <doi.org/10.5964/bioling.13153>.

Kawahara, Shigeto; Noto, Atsushi & Kumagai, Gakuji 2018. Sound symbolic patterns in Pokémon names. *Phonetica* 75,3. 219-244. <DOI: 10.1159/000484938>.

Kayne, Richard S. 1994. *The antisymmetry of syntax*. Cambridge, MA: MIT Press.

Keine, Stefan 2020. *Probes and their horizons*. Cambridge, MA: MIT Press.

Kempson, Ruth; Meyer Viol, Wilfried & Gabbay, Dov M. 2001. *Dynamic Syntax: The Flow of Language Understanding*. Wiley.

Kennedy, Christopher 2015. A “de-Fregean” semantics (and neo-Gricean pragmatics) for modified and unmodified numerals. *Semantics & Pragmatics* 8. 1-44. <dx.doi.org/10.3765/sp.8.1>.

Kerr, Dara 2024. How Memphis became a battleground over Elon Musk’s xAI supercomputer. *NPR* 11 September 2024. <www.npr.org/2024/09/11/6588134/elon-musk-ai-xai-supercomputer-memphis-pollution>.

Kharitonov, Eugene & Chaabouni, Rahma 2021. What they do when in doubt: A study of inductive biases in seq2seq learners. In *ICLR 2021 Conference Track*. Online: OpenReview.

Kim, Najoung; Patel, Roma; Poliak, Adam; Wang, Alex; Xia, Patrick; McCoy, R. Thomas; Tenney, Ian; Ross, Alexis; Linzen, Tal & van Durme, Benjamin 2019. Probing what different NLP tasks teach machines about function word comprehension. <[arXiv:1904.11544](https://arxiv.org/abs/1904.11544)>.

Kingma, Diederik P. & Ba, Jimmy Lei 2015. Adam: A method for stochastic optimization. In *ICLR 2015 Conference Track*. San Diego, CA: OpenReview.

Kirov, Christo & Cotterell, Ryan 2018. Recurrent Neural Networks in Linguistic Theory: Revisiting Pinker and Prince (1988) and the Past Tense Debate. *Transactions of the Association for Computational Linguistics* 6 (December). 651-665. <doi.org/10.1162/tacl_a_00247>.

Kitaev, Nikita; Cao, Steven & Klein, Daniel 2019. Multilingual constituency parsing with self-attention and pre-training. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (ACL 2019). 3499-3505.

Kitchin, Rob 2014. Big Data, new epistemologies and paradigm shifts. *Big Data & Society* 1,1. <DOI: 10.1177/2053951714528481>.

Klein, Daniel & Manning, Christopher D. 2003. Accurate unlexicalized parsing. In *Proceedings of the 41st Meeting of the Association for Computational Linguistics*. 423-430.

Kleyko, Denis; Rachkovskij, Dmitri; Osipov, Evgeny & Rahimi, Abbas

2023. A survey on hyperdimensional computing aka vector symbolic architectures, parts 1 and 2. *ACM Computing Surveys* 55. 130. <doi.org/10.1145/3538531>.

Klimova, Blanka; Pikhart, Marcel & Al-Obaydi, Liqaa Habeb 2024. Exploring the potential of ChatGPT for foreign language education at the university level. *Frontiers in Psychology* 15. <DOI: 10.3389/fpsyg.2024.1269319>.

Knight, Chris 2016. *Decoding Chomsky: Science and Revolutionary Politics*. New Haven, CT: Yale University Press. <DOI: 10.12987/9780300222159>.

Kobele, Gregory M. 2023. Minimalist Grammars and Decomposition. In Kleanthes, Grohmann & Leivada, Evelina (eds.), *The Cambridge Handbook of Minimalism*. Cambridge University Press.

Kobele, Gregory M.; Gerth, Sabrina & Hale, John T. 2013. Memory resource allocation in top-down Minimalist parsing. In Morrill, Glyn & Nederhof, Mark-Jan (eds.), *Formal grammar: 17th and 18th international conferences, FG 2012, Opole, Poland, August 2012, Revised selected papers, FG 2013, Düsseldorf, Germany, August 2013*. 32-51. Berlin / Heidelberg: Springer. <doi.org/10.1007/978-3-642-39998-5_3>.

Kodner, Jordan; Payne, Sarah & Heinz, Jeffrey 2023. Why linguistics will thrive in the 21st century: A reply to Piantadosi (2023). <arxiv.org/abs/2308.03228>.

Koerner, Konrad 1983. The Chomskyan 'revolution' and its historiography: A few critical remarks. *Language & Communication* 3,2. 147-169. <DOI: 10.1016/0271-5309(83)90012-5>.

Kojima, Takeshi; Gu, Shixiang (Shane); Reid, Machel; Matsuo, Yutaka & Iwasawa, Yusuke 2022. Large language models are zero-shot reasoners. In Koyejo, S.; Mohamed, S.; Agarwal, Al; Belgrave, D.; Cho, K. & Oh, A. (eds.), *Advances in Neural Information Processing Systems 35 (NeurIPS 2022) Main Conference Track*. New Orleans, LA: Curran Associates, Inc. 22199-22213.

Kolmogorov, Andrey N. 1963. On Tables of Random Numbers. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)* 25,4. 369-376.

Kuhn, Thomas 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.

Kwon, Diana 2024. AI is complicating plagiarism. How should scientists respond? *Nature*. <DOI: 10.1038/d41586-024-02371-z>.

Lake, Brenden M. & Baroni, Marco 2023. Human-like systematic generalization through a meta-learning neural network. *Nature* 623. 115-121. <doi.org/10.1038/s41586-023-06668-3>.

Lakretz, Yair; Hupkes, Dieuwke; Vergallito, Alessandra; Marelli, Marco; Baroni, Marco & Dehaene, Stanislas 2021. Mechanisms for handling nested dependencies in neural-network language models and humans. *Cognition* 213. 1-24. <DOI: 10.1016/j.cognition.2021.104699>. <www.sciencedirect.com/science/article/pii/S0010027721001189>.

Lakretz, Yair; Kruszewski, German; Desbordes, Theo; Hupkes, Dieuwke; Dehaene, Stanislas & Baroni, Marco 2019. The emergence of number and syntax units in LSTM language models. In Burstein, Jill; Doran, Christy & Solorio, Thamar (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Vol. 1*. Minneapolis, MN: Association for Computational Linguistics. 11-20. <DOI: 10.18653/v1/N19-1002>.

Lan, Nur; Chemla, Emmanuel & Katzir, Roni 2024. Large language models and the argument from the poverty of the stimulus. *Linguistic Inquiry*. 1-28. <doi.org/10.1162/ling_a_00533>.

Lan, Nur; Geyer, Michal; Chemla, Emmanuel & Katzir, Roni 2022. Minimum Description Length Recurrent Neural Networks. *Transactions of the Association for Computational Linguistics* 10 (July). 785-99. <doi.org/10.1162/tacl_a_00489>.

Landman, Fred 2003. Predicate-argument mismatches and the adjectival theory of indefinites. In Coene, M. & d'Hulst, Y. (eds.), *From NP to DP: The syntax and semantics of noun phrases*. Volume 1. 211-237. Amsterdam: John Benjamins.

Lasnik, Howard & Lidz, Jeffrey L. 2016. The argument from the poverty of the stimulus. In Roberts, Ian (ed.), *The Oxford Handbook of Universal Grammar*. Oxford: Oxford University Press. 221-248.

Latour, Bruno 1984. *Les Microbes: Guerre et paix, suivi de Irréductions*. Paris, France: A. M. Métailié.

Law, John & Lodge, Peter 1984. *Science for Social Scientists*. London: Palgrave Macmillan UK. <DOI: 10.1007/978-1-349-17536-9>.

Lawson, Alex 2024. Google to buy nuclear power for AI datacentres in 'world first' deal. *Guardian* 15 October 2024. <www.theguardian.com/technology/2024/oct/15/google-buy-nuclear-power-ai-datacentres-kairos-power>.

Lee, So Young & De Santo, Aniello. A computational view into the structure of attachment ambiguities in Chinese and Korean. In *Proceedings of the north east linguistics society*. 189-198.

Levesque, Hector J. 2014. On our best behaviour. *Artificial Intelligence* 212. 27-35. <doi.org/10.1016/j.artint.2014.03.007>.

Levshina, Natalia 2023. *Communicative efficiency: Language structure and use*. Cambridge: Cambridge University Press.

Levy, Roger 2008. Expectation-based syntactic comprehension. *Cognition* 106,3. 1126-1177.

Li, Jixing; Bhattacharji, Shohini; Zhang, Shulin; Franzluebbers, Berta; Luh, Wen-Ming; Spreng, R. Nathan; Brennan, Jonathan R.; Yang, Yiming; Pallier, Christophe & Hale, John 2022. *Le Petit Prince* multilingual naturalistic fMRI corpus. *Scientific Data* 9. 530. <doi.org/10.1038/s41597-022-01625-7>.

Li, Jixing & Hale, John 2019. Grammatical predictors for fMRI time-courses.

In Berwick, Robert C. & Stabler, Edward P. (eds.), *Minimalist Parsing*. Oxford, UK: Oxford University Press. 159-173. <doi.org/10.1093/oso/9780198795087.003.0007>.

Li, Ming & Vitányi, Paul 2008. *An Introduction to Kolmogorov Complexity and Its Applications*. New York: Springer. <doi.org/10.1007/978-0-387-49820-1>.

Lidz, Jeffrey & Gleitman, Lila R. 2004. Argument structure and the child's contribution to language learning. *Trends in Cognitive Sciences* 8,4.

Lillicrap, Timothy P.; Santoro, Adam; Marris, Luke; Akerman, Colin J. & Hinton, Geoffrey 2020. Backpropagation and the Brain. *Nature Reviews Neuroscience* 21,6. 335-46. <doi.org/10.1038/s41583-020-0277-3>.

Lin, Stephanie; Hilton, Jacob & Evans, Owain 2022. TruthfulQA: Measuring how models mimic human falsehoods. In Muresan, Smaranda; Nakov, Preslav & Villavicencio, Aline (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. Vol. 1. Dublin, Ireland: Association for Computational Linguistics. 3214-3252. <DOI: 10.18653/v1/2022.acl-long.229>.

Ling, Jacqueline 2001. Power of a human brain. In *Physics Factbook*. <hypertextbook.com/facts/2001/JacquelineLing.shtml>.

Link, Godehard 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In Bauerle, Rainer; Schwarze, Christoph & von Stechow, Arnim (eds.), *Meaning, Use, and the Interpretation of Language*. Berlin / New York: de Gruyter. 302-323.

Linzen, Tal & Baroni, Marco 2021. Syntactic structure from deep learning. *Annual Review of Linguistics* 7. 195-212. <DOI: 10.1146/annurev-linguistics-032020-051035>.

Linzen, Tal; Dupoux, Emmanuel & Goldberg, Yoav 2016. Assessing the Ability of LSTMs to Learn Syntax-Sensitive Dependencies. *Transactions of the Association for Computational Linguistics* 4. 521-535. <doi.org/10.1162/tacl_a_00115>.

Liu, Lei 2023. Processing advantages of end-weight. *Proceedings of the Society for Computation in Linguistics* 6. 250-258.

Lohninger, Magdalena & Wurmbrand, Susi 2025. Typology of Complement Clauses. In Benz, Anton; Frey, Werner; Gärtner, Hans-Martin; Krifka, Manfred; Schenner, Mathias & Źygis, Marzena (eds.), *Handbook of clausal embedding*. Berlin: Language Science Press.

Longobardi, Giuseppe 1994. Reference and proper names: A theory of N-movement in syntax and logical form. *Linguistic Inquiry* 25. 609-665.

Manning, Christopher D.; Clark, Kevin; Hewitt, John; Khandelwal, Uravashi & Levy, Omer 2020. Emergent linguistic structure in artificial neural networks trained by self-supervision. In Gavish, Matan (ed.), *Proceedings of the National Academy of Science of the United States of America* 117. 30046-30054. <DOI: 10.1073/pnas.1907367117>.

Manzini, Maria Rita 1983. Syntactic conditions on phonological rules. *MIT Working Papers in Linguistics* 5. 1-9.

Marantz, Alec 2019. What do linguists do? In *The Julius Silver, Roslyn S. Silver, and Enid Silver Winslow Dialogues in Arts and Science*, New York University. <as.nyu.edu/content/dam/nyu-as/as/documents/silverdialogues/SilverDialogues_Marantz.pdf>.

Marcus, Gary 2022. Noam Chomsky and GPT-3 [Blog post]. *Marcus on AI*. <garymarcus.substack.com/p/noam-chomsky-and-gpt-3>. Last accessed 24/02/2025.

Marcus, Mitchell *et al.* 1994. The Penn Treebank: Annotating predicate argument structure. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey*.

Marr, David 1982. *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA: Freeman.

Marr, David & Poggio, Tomaso 1976. *From Understanding Computation to Understanding Neural Circuitry*. Cambridge, MA: MIT Press.

Martinetz, Julius; Linse, Christoph & Martinetz, Thomas 2024. Rethinking generalization of classifiers in separable classes scenarios and over-parameterized regimes. *International Joint Conference on Neural Networks 2024*. 1-10. <doi.org/10.1109/IJCNN60899.2024.10650680>.

Marvin, Rebecca & Linzen, Tal 2018. Targeted Syntactic Evaluation of Language Models. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 1192-1202. <doi.org/10.18653/v1/D18-1151>.

May, Robert 1985. *Logical form: Its structure and derivation* (Vol. 12). Cambridge, MA: MIT Press.

Mayer, Connor & Major, Travis 2018. A challenge for tier-based strict locality from Uyghur backness harmony. In Foret, Annie; Kobelev, Greg & Pogodalla, Sylvain (eds.), *Proceedings of formal grammar 2018*. Berlin: Springer. 62-83.

McCawley, James D. 1976. Introduction. In McCawley, James D. (ed.), *Notes From the Linguistic Underground*. New York, NY: Academic Press. 1-19.

McClelland, James L. & Rumelhart, David E. 1991. *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs, and Exercises*. 2nd print. Computational Models of Cognition and Perception. Cambridge, MA: MIT Press.

McCoy, Richard; Frank, Robert & Linzen, Tal 2018. Revisiting the poverty of the stimulus: Hierarchical generalization without a hierarchical bias in recurrent neural networks. In *Proceedings of the Annual Meeting of the Cognitive Science Society*. Madison, WI: Cognitive Science Society. 2096-2101.

McCoy, R. Thomas; Yao, Shunyu; Friedman, Dan; Hardy, Matthew & Griffiths, Thomas L. 2023. Embers of autoregression: Understanding large language models through the problem they are trained to solve. <arxiv.org/abs/2309.13638>.

McCullough, Gretchen 2019. *Because Internet: Understanding the New Rules of Language*. New York, NY: Riverhead Books.

McGee, Thomas & Blank, Idan 2024. Evidence against syntactic encapsulation in large language models. *Procs. Cognitive Science Society* 46.

McKenzie, Ian R.; Lyzhov, Alexander; Pieler, Michael Martin; Parrish, Alicia; Mueller, Aaron; Prabhu, Ameya; McLean, Euan; Shen, Xudong; Cavanagh, Joe, Gritsevskiy, Andrew George *et al.* 2023. Inverse scaling: When bigger isn't better. *Transactions on Machine Learning Research*.

McNally, Louise & Boleda, Gemma 2004. Relational adjectives as properties of kinds. *Empirical Issues in Syntax and Semantics* 5. 179-196. <doi.org/ISSN1769-7158>.

Merrill, William; Sabharwal, Ashish & Smith, Noah A. 2022. Saturated transformers are constant-depth threshold circuits. *Transactions of the Association for Computational Linguistics* 10. 843-856. <DOI: 10.1162/tacl_a_00493>.

Michaelis, Jens 2001. Derivational Minimalism Is Mildly Context-Sensitive. In Moortgat, Michael (ed.), *Logical Aspects of Computational Linguistics* (Vol. 2014). Berlin / Heidelberg: Springer. 179-198. <doi.org/10.1007/3-540-45738-0_11>.

Mikolov, Tomáš 2012. *Statistical Language Models Based on Neural Networks*. PhD dissertation. Brno University of Technology, Brno, Czech Republic.

Milewski, Bartosz 2020. *Category Theory for Programmers*. <bartoszmilewski.com>.

Miller, George A. & Chomsky, Noam 1963. Finitary Models of Language Users. In Luce, D. (ed.), *Handbook of Mathematical Psychology*. John Wiley & Sons. 2-419.

Milway, Daniel 2023. A response to Piantadosi (2023). <lingbuzz/007264>.

Mishra, Swaroop; Khashabi, Daniel; Baral, Chitta & Hajishirzi, Hannaneh 2022. Cross-task generalization via natural language crowdsourcing instructions. In Muresan, Smaranda; Nakov, Preslav & Villavicencio, Aline (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. Vol. 1. Dublin, Ireland: Association for Computational Linguistics. 3470-3487. <DOI: 10.18653/v1/2022.acl-long.244>.

Mitchell, J.; Kazanina, Nina; Houghton, Conor J.; Bowers, Jeffrey S. 2019. Do LSTMs know about Principle C? In *2019 Conference on Cognitive Computational Neuroscience*.

Mollica, Frank & Piantadosi, Stephen 2019. Humans store about 1.5 megabytes of information during language acquisition. *Royal Society Open Science* 6,3.

Mollica, Frank & Piantadosi, Stephen 2022. Meaning without reference in large language models. <arXiv:2208.02957>.

Momma, Shota & Phillips, Colin 2018. The Relationship Between Parsing and Generation. *Annual Review of Linguistics* 4,1. 233-254. <doi.org/10.1146/annurev-linguistics-011817-045719>.

Moro, Andrea 2023. Embodied syntax: Impossible languages and the irreducible difference between humans and machines. *Sistemi intelligenti* 2.

321-328. <doi.org/10.1422/108132>.

Moro, Andrea; Greco, Matteo & Cappa, Stefano F. 2023. Large languages, impossible languages and human brains. *Cortex* 167. 82-85. <doi.org/10.1016/j.cortex.2023.07.003>.

Müller, Stefan 2024. Large language models: The best linguistic theory, a wrong linguistic theory, or no linguistic theory at all. *Zeitschrift für Sprachwissenschaft*.

Mullins, Nicholas C. 1975. A sociological theory of scientific revolution. In Knorr, Karin D.; Strasser, Hermann & Zilian, Hans Georg (eds.), *Determinants and Controls of Scientific Development*. Dordrecht, Netherlands: Springer Netherlands. 185-203.

Murray, Stephen O. 1994. *Theory Groups and the Study of Language in North America*. Amsterdam, Netherlands: John Benjamins.

Murty, Shikhar; Sharma, Pratyusha; Andreas, Jacob & Manning, Christopher D. 2022. Characterizing intrinsic compositionality in transformers with tree projections.

Naveed, Humza; Asad Ullah Khan; Shi Qiu; Saqib, Muhammad; Anwar, Saeed; Usman, Muhammad; Akhtar, Naveed; Barnes, Nick & Mian, Ajmal 2024. A comprehensive overview of large language models. <arxiv.org/abs/2307.06435>.

Newmeyer, Frederick J. 1980. *Linguistic theory in America: The first quarter century of Transformational Generative Grammar*. New York: Academic Press.

Newmeyer, Frederick J. 1986. Has there been a 'Chomskyan revolution' in linguistics? *Language* 62,1. 1-18. <DOI: 10.2307/415597>.

Newmeyer, Frederick J. 2004. Against a parameter-setting approach to typological variation. *Linguistic Variation Yearbook* 4,1. 181-234. <[doi:10.1075/livy.4.06new](https://doi.org/10.1075/livy.4.06new)>.

Newmeyer, Frederick J. 2021. Complexity and relative complexity in generative grammar. *Frontiers in Communication* 6. <[doi:10.3389/fcomm.2021.614352](https://doi.org/10.3389/fcomm.2021.614352)>.

Newmeyer, Frederick J. & Emonds, Joseph 1971. The linguist in American society. In *Papers from the Seventh Regional Meeting of the Chicago Linguistic Society*. Chicago, IL: Chicago Linguistic Society. 285-303.

Nivre, Joakim; Agić, Željko; Ahrenberg, Lars; Antonsen, Lene; Aranzabe, María Jesus; Asahara, Masayuki; Ateyah, Luma; Attia, M.; Atutxa, A.; Augustinus, L. et al. 2017. *Universal Dependencies 2.1*.

Norvig, Peter 2017. On Chomsky and the two cultures of statistical learning. In Pietsch, Wolfgang; Wernecke, Jörg & Ott, Maximilian (eds.), *Berechenbarkeit der Welt? Philosophie und Wissenschaft im Zeitalter von Big Data*. Wiesbaden, Germany: Springer Fachmedien. 61-83.

Nosengo, Nicola 2014. *I robot ci guardano: Aerei senza pilota, chirurghi a distanza e automi solidali*. Bologna: Zanichelli.

Noy, Shakked & Zhang, Whitney 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381, 6654. 187-192. <DOI: 10.1126/science.adh2586>.

Nvidia n.d. *meta/llama-3.1-405b-instruct*. *Nvidia API reference*. <docs.api.nvidia.com/nim/reference/meta-llama-3_1-405b>.

Nye, Maxwell; Andreassen, Anders Johan; Gur-Ari, Guy; Michalewski, Henryk; Austin, Jacob; Bieber, David; Dohan, David; Lewkowycz, Aitor; Bosma, Maarten; Luan, David; Sutton, Charles & Odena, Augustus (2022). Show your work: Scratchpads for intermediate computation with language models. In *ICLR 2022 Workshop DL4C*. Online: OpenReview.

Oepen, Stephan; Toutanova, Kristina; Shieber, Stuart; Manning, Christopher; Flickinger, Dan & Brants, Thorsten 2022. The LinGO Redwoods treebank: Motivation and preliminary applications. In *COLING 2002: The 17th International Conference on Computational Linguistics: Project Notes*. Taipei, Taiwan: Association for Computational Linguistics.

Oerter, Robert 2006. *The theory of almost everything: The Standard Model, the unsung triumph of modern physics*. New York: Pi Press.

Oh, Byung-Doh & Schuler, William 2023. Why does surprisal from larger transformer-based language models provide a poorer fit to human reading times? *Transactions of the Association for Computational Linguistics* 11. 336-350. <DOI: 10.1162/tacl_a_00548>.

OpenAI 2023. *GPT-4 Technical Report* <arxiv.org/abs/2303.08774>.

Ouyang, Long; Wu, Jeff; Jiang, Xu; Almeida, Diogo; Wainwright, Carroll L.; Mishkin, Pamela; Zhang, Chong; Agarwal, Sandhini; Slama, Katarina; Ray, Alex *et al.* 2022. Training language models to follow instructions with human feedback. <[arXiv:2203.02155](https://arxiv.org/abs/2203.02155)>.

Ozaki, Satoru; Santo, Aniello De; Linzen, Tal & Dillon, Brian 2024. CCG parsing effort and surprisal jointly predict RT but underpredict garden-path effects. *Society for Computation in Linguistics* 7. 362-364. <doi.org/10.7275/scil.2229>.

Papineni, Kishore; Roukos, Salim; Ward, Todd & Zhu, Wei-Jing 2001. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02*, 311. Philadelphia, Pennsylvania: Association for Computational Linguistics. <doi.org/10.3115/1073083.1073135>.

Park, Peter S.; Goldstein, Simon; O’Gara, Aidan; Chen, Michael & Hendrycks, Dan 2024. AI deception: A survey of examples, risks, and potential solutions. *Patterns* 5, 5. 100988. <DOI: 10.1016/j.pattern.2024.100988>.

Pascanu, Razvan; Mikolov, Tomas & Bengio, Yoshua 2013. On the difficulty of training recurrent neural networks. In Dasgupta, Sanjoy & McAllester, David (eds.), *ICML’13: Proceedings of the 30th International Conference on International Conference on Machine Learning*. Vol. 28. Atlanta, GA: Proceedings of Machine Learning Research. 1310-1318.

Pasternak, Robert & Graf, Thomas 2021. Cyclic scope and processing difficulty in a Minimalist parser. *Glossa* 6. 1-34. <doi.org/10.5334/gjgl.1209>.

Pasteur, Louis 1876. *Études sur la bière, ses maladies, causes qui les provoquent, procédé pour la rendre inaltérable, avec une théorie nouvelle de la fermentation*. Paris: Gauthier-Villars.

Pasteur, Louis 1880. De l'extension de la théorie des germes à l'étiologie de quelques maladies communes. In *Comptes rendus hebdomadaires des séances de l'Académie des sciences*. Vol. 90. Paris: Gauthier-Villars. 1033-1034.

Pasteur, Louis; Joubert, Jules & Chamberland, Charles 1878. La théorie des germes et ses applications à la médecine et à la chirurgie. In *Comptes rendus hebdomadaires des séances de l'Académie des sciences*. Vol. 86. Paris: Gauthier-Villars. 1037-1043.

Pater, Joe 2019. Generative linguistics and neural networks at 60: Foundation, friction, and fusion. *Language* 95. 41-74. <doi.org/10.1353/lan.2019.0009>.

Pearl, Lisa 2022. Poverty of the stimulus without tears. *Language Learning and Development* 18,4. 415-454. <DOI: 10.1080/15475441.2021.1981908>.

Pennington, Jeffrey; Socher, Richard & Manning, Christopher D. 2014. Glove: Global vectors for word representation. *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532-1543.

Pereira, Fernando 2000. Formal grammar and information theory: Together again? *Philosophical Transactions: Mathematical, Physical and Engineering Sciences* 358,1769. 1239-1253.

Perez, Ethan; Huang, Saffron; Song, Francis; Cai, Trevor; Ring, Roman; Aslanides, John; Glaese, Amelia; McAleese, Nat & Irving, Geoffrey 2022. Red teaming language models with language models. In Goldberg, Yoav; Kozareva, Zornitsa & Zhang, Yue (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics. 3419-3448. <DOI: 10.18653/v1/2022.emnlp-main.225>.

Pesetsky, David 2024. Is there an LLM challenge for Linguistics? Or is there a Linguistics challenge for LLMs?. Paper presented at the Academia Română, Bucarest, 22 May 2024.

Petroni, Fabio; Rocktäschel, Tim; Riedel, Sebastian; Lewis, Patrick; Bakhtin, Anton; Wu, Yuxiang & Miller, Alexander 2019. Language models as knowledge bases? In Inui, Kentaro; Jiang, Jing; Ng, Vincent & Wan, Xiaojun (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics. 2463-2473. <DOI: 10.18653/v1/D19-1250>.

Phillips, Colin 1996. *Order and structure*. PhD dissertation. Cambridge, MA: MIT Press.

Phillips, Colin 2003. Linear order and constituency. *Linguistic Inquiry* 34. 37-90.

Piantadosi, Steven T. 2023. Modern language models refute Chomsky's approach to language. <lingbuzz.net/lingbuzz/007180>.

Piantadosi, Steven T. 2024. Modern language models refute Chomsky's approach to language. In Gibson, Edward & Poliak, Moshe (eds.), *From*

fieldwork to linguistic theory: A tribute to Dan Everett. Berlin: Language Science Press. 353-414.

Pinker, Steven 1984. *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.

Plate, Tony A. 1994. *Holographic Reduced Representation*. Stanford: CSLI.

Plato 380 BCE. *Meno*.

Poggio, Thomas; Rifkin, Ryan; Niyogi, Partha & Mukherjee, Sayan 2004. General conditions for predictivity in learning theory. *Nature* 428. 419-422. <doi.org/10.1038/nature02341>.

Pollard, Carl & Sag, Ivan A. 1994. *Head-Driven Phrase Structure Grammar*. Chicago, IL: University of Chicago Press.

Pollock, Jean Yves 1989. Verb movement, universal grammar, and the structure of IP. *Linguistic inquiry* 20.3. 365-424.

Popper, Karl 1934. *Logik der Forschung*. Berlin: Springer. <doi.org/10.1007/978-3-7091-4177-9>.

Prasanna, Sai; Rogers, Anna & Rumshisky, Anna 2020. When BERT plays the lottery, all tickets are winning. In Webber, Bonnie; Cohn, Trevor; He, Yulan & Liu, Yang (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics. 3208-3229. <DOI: 10.18653/v1/2020.emnlp-main.259>.

Pullum, Geoffrey K. & Scholz, Barbara C. 2002. Empirical assessment of stimulus poverty arguments. *The Linguistic Review* 18,1-2. 9. <DOI: 10.1515/tlir.19.1-2.9>.

Purnell, Thomas; Idsardi, William & Baugh, John 1999. Perceptual and phonetic experiments on American English dialect identification. *Journal of Language and Social Psychology* 18,1. 10-30. <DOI: 10.1177/0261927X99018001002>.

Quine, Willard Van Orman 1960. *Word and Object*. Cambridge, MA: MIT Press. <doi.org/10.7551/mitpress/9636.001.0001>.

Quinlan, Philip T. (ed.) 2004. *Connectionist Models of Development* (0 ed.). Psychology Press. <doi.org/10.4324/9780203494028>.

Radford, Alec; Narasimhan, Karthik; Salimans, Tim; Sutskever, Ilya *et al.* 2018. *Improving language understanding by generative pre-training*.

Radford, Alec; Wu, Jeffrey; Amodei, Dario; Clark, Jack; Brundage, Miles & Sutskever, Ilya 2019a. Better language models and their implications [Blog post]. *OpenAI Research*. <openai.com/index/better-language-models>. Last accessed 24/02/2025.

Radford, Alec; Wu, Jeffrey; Child, Rewon; Luan, David; Amodei, Dario & Sutskever, Ilya 2019b. *Language Models Are Unsupervised Multitask Learners*. Technical report. San Francisco, CA: OpenAI.

Radford, Andrew 1997. *Syntax: A Minimalist Introduction*. Cambridge: Cambridge University Press.

Radford, Andrew 2016. *Analysing English Sentences, Second Edition*. Cambridge: Cambridge University Press.

Rafailov, Rafael; Sharma, Archit; Mitchell, Eric; Ermon, Stefano; Manning, Christopher D. & Finn, Chelsea 2023. Direct preference optimization: Your language model is secretly a reward model. In *ICLR 2023 Conference Track*. Kigali, Rwanda: OpenReview.

Raman, Raghu *et al.* 2024. Fake news research trends, linkages to generative artificial intelligence and sustainable development goals. *Helion* e24727. <DOI: 10.1016/j.heliyon.2024.e24727>.

Rasin, Ezer; Berger, Iddo; Lan, Nur; Shefi, Itamar & Katzir, Roni 2021. Approaching explanatory adequacy in phonology using minimum description length. *Journal of Language Modelling* 9,1. 17-66. <doi.org/10.15398/jlm.v9i1.266>.

Rawski, Jonathan & Heinz, Jeffrey 2019. No free lunch in linguistics or machine learning: Response to Pater. *Language* 95. 125-135.

Raymond, Louise & O'Reilly, Tim 1999. *The Cathedral and the Bazaar* (1st ed.). USA: O'Reilly & Associates, Inc.

Reinhart, Tanya 1976. *The syntactic domain of anaphora*. Cambridge, MA: MIT Press.

Retoré, Christian (ed.), *Logical Aspects of Computational Linguistics: Lecture Notes in Computer Science*. Berlin: Springer. 68-95.

Rickford, John R. & King, Sharese 2016. Language and linguistics on trial: Hearing Rachel Jeantel (and other vernacular speakers) in the courtroom and beyond. *Language* 92,4. 948-988.

Riesenhuber, Maximilian & Poggio, Tomaso 1999. Hierarchical models of object recognition in cortex. *Nature Neuroscience* 2,11. 1019-1025. <doi.org/10.1038/14819>.

Riezler, Stefan; Holloway King, Tracy; Kaplan, Ronald M.; Crouch, Richard; Maxwell, John T. III & Johnson, Mark 2002. Parsing the Wall Street Journal using a Lexical-Functional Grammar and discriminative estimation techniques. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Philadelphia: Association for Computational Linguistics. 271-278.

Rissanen, Jorma 1978. Modeling by shortest data description. *Automatica* 14,5. 465-471. <doi.org/10.1016/0005-1098(78)90005-5>.

Rissanen, Jorma 1987. Stochastic Complexity. *Journal of the Royal Statistical Society: Series B (Methodological)* 49,3. 223-239. <doi.org/10.1111/j.2517-6161.1987.tb01694.x>.

Ritter, Elizabeth & Wiltschko, Martina 2014. The composition of INFL. An exploration of tense, tenseless languages and tenseless constructions. *Natural Language and Linguistic Theory* 32. 1331-1386.

Ritter, Elizabeth 1991. Two functional categories in Noun Phrases: Evidence from Modern Hebrew. *Syntax and Semantics* 25.

Rizzi, Luigi 1990. *Relativized minimality*. Cambridge, MA: MIT Press.

Rizzi, Luigi 1997. The Fine Structure of the Left Periphery. In Haegeman, Liliane (ed.), *Elements of Grammar*. Dordrecht: Springer Netherlands. 281-337. <doi.org/10.1007/978-94-011-5420-8_7>.

Rizzi, Luigi (ed.) 2004. *The structure of CP and IP*. Oxford, UK: Oxford University Press.

Rizzi, Luigi 2013. Locality. *Lingua* 130. 169-186.

Rizzi, Luigi 2016. Labeling, maximality and the head-phrase distinction. *The Linguistic Review* 33.1. 103-127.

Rizzi, Luigi 2021. *Complexité des structures linguistiques, simplicité des mécanismes du langage*, Leçon inaugurale, 2021, Collège de France – Fayard, Paris. English translation: *Complexity of Linguistic Structures, Simplicity of Language Mechanisms* (2024). OpenEdition Books, Collège de France. <DOI: 10.4000/books.cdf.16006>.

Rizzi, Luigi & Cinque, Guglielmo 2016. Functional Categories and Syntactic Theory. *Annual Review of Linguistics* 2.1. 139-163. <doi.org/10.1146/annurev-linguistics-011415-040827>.

Rizzi, Luigi & Savoia, Leonardo 1993. Conditions on /u/ propagation in Southern Italian Dialects: A Locality Parameter for Phonosyntactic Processes. In Belletti, A. (ed.), *Syntactic Theory and the Dialects of Italy*. Turin: Rosenberg & Sellier.

Roberts, Ian 2017. The final-over-final condition in DP: Universal 20 and the nature of demonstratives. In Sheehan, Michelle; Biberauer, Theresa; Roberts, Ian & Holmberg, Anders (eds.), *The Final-over-Final Condition: A Syntactic Universal* (Vol. 76). Cambridge, MA: MIT Press. 151.

Roberts, Ian 2019. *Parameter Hierarchies and Universal Grammar* (1st ed.). Oxford, UK: Oxford University Press. <doi.org/10.1093/oso/9780198804635.001.0001>.

Rogers, Anna; Kovaleva, Olga & Rumshisky, Anna 2021. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics* 8. 842-866.

Ross, John Robert 1967. *Constraints on variables in syntax*. Cambridge, MA: MIT Press.

Rumelhart, David E.; Hinton, Geoffrey E. & Williams, Ronald J. 1986. Learning representations by back-propagating errors. *Nature* 323,6088. 533-536. <DOI: 10.1038/323533a0>.

Rumelhart, David E. & McClelland, James L. 1986. On Learning the Past Tenses of English Verbs. In *Parallel Distributed Processing*. Cambridge, MA: MIT Press. <doi.org/10.7551/mitpress/5237.003.0008>.

Rumelhart, David E.; McClelland, James L. & PDP Research Group (eds.) 1999. *Parallel distributed processing. 1: Foundations*. 12th print. Cambridge, MA: MIT Press. <doi.org/10.7551/mitpress/5236.001.0001>.

Russell, Bertrand 1947. *Human Knowledge: Its Scope and Limits*. New York, NY: Simon and Schuster.

Sampson, Geoffrey 1997. *Educating Eve: The ‘language instinct’ debate*. London / Washington, DC: Cassell.

Sanh, Victor; Webson, Albert; Raffel, Colin; Bach, Stephen; Sutawika, Lintang; Alyafeai, Zaid; Chaffin, Antoine; Stiegler, Arnaud; Raja, Arun;

Dey, Manan *et al.* 2022. *Multitask prompted training enables zero-shot task generalization*. In *ICLR 2022 Conference Track*. Online: OpenReview.

Sarlin, Paul-Edouard; DeTone, Daniel; Malisiewicz, Tomasz & Rabinovich, Andrew 2020. Superglue: Learning feature matching with graph neural networks. <arxiv.org/abs/1911.11763>.

Sartran, Laurent; Barrett, Samuel; Kuncoro, Adhiguna; Stanojević, Miloš; Blunsom, Phil & Dyer, Chris 2022. Transformer Grammars: Augmenting Transformer Language Models with Syntactic Inductive Biases at Scale. *Transactions of the Association for Computational Linguistics* 10 (December). 1423-39. <doi.org/10.1162/tacl_a_00526>.

Sathish, Vishwas; Lin, Hannah; Kamath, Aditya K. & Nyayachavadi, Anish 2024. LLeMpower: Understanding disparities in the control and access of large language models. <[arXiv:2404.09356](https://arxiv.org/abs/2404.09356)>.

Savitch, Walter J. 1993. Why it might pay to assume that languages are infinite. *Annals of Mathematics and Artificial Intelligence* 8. 17-25.

Shannon, C. E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27,3. 379-423. <DOI: [10.1002/j.1538-7305.1948.tb01338.x](https://doi.org/10.1002/j.1538-7305.1948.tb01338.x)>.

Shieber, Stuart M. 1985. Evidence against the Context-Freeness of Natural Language. *Linguistics and Philosophy* 8,3. 333-43. <doi.org/10.1007/BF00630917>.

Siegelman, Noam; Schroeder, Sascha; Acartürk, Cengiz; Ahn, Hee-Don; Alexeeva, Svetlana; Amenta, Simona; Bertram, Raymond; Bonandolini, R.; Brysbaert, M.; Chernova, D.; Da Fonseca, S. M.; Dirix, N.; Duyck, W.; Fella, A.; Frost, R.; Gattei, C. A.; Kalaitzi, A.; Kwon, N.; Lõo, K.; ... Kuperman, V. 2022. Expanding horizons of cross-linguistic research on reading: The Multilingual Eye-movement Corpus (MECO). *Behavior Research Methods* 54,6. 2843-2863. <doi.org/10.3758/s13428-021-01772-6>.

Smith, Nathaniel J. & Levy, Roger 2013. The effect of word predictability on reading time is logarithmic. *Cognition* 128,3. 302-319. <DOI: [10.1016/j.cognition.2013.02.013](https://doi.org/10.1016/j.cognition.2013.02.013)>.

Smolensky, Paul 1990. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence* 46. 159-216. <[doi.org/10.1016/0004-3702\(90\)90007-m](https://doi.org/10.1016/0004-3702(90)90007-m)>.

Solomonoff, Ray J. 1960. *A Preliminary Report on a General Theory of Inductive Inference*. United States Air Force, Office of Scientific Research. <books.google.it/books?id=SuTHtgAACAAJ>.

Spitale, Giovanni; Biller-Andorno, Nikola & Germani, Federico 2023. AI model GPT-3 (dis)informs us better than humans. *Science Advances* 9, 26. <DOI: [10.1126/sciadv.adh1850](https://doi.org/10.1126/sciadv.adh1850)>.

Sprouse, Jon & Almeida, Diogo 2017. Design sensitivity and statistical power in acceptability judgment experiments. *Glossa* 2,1. 1-32. <doi.org/10.5334/gjgl.236>.

Sprouse, Jon & Hornstein, Norbert (eds.) 2013. *Experimental Syntax and*

Island Effects (1st ed.). Cambridge University Press. <doi.org/10.1017/CBO9781139035309>.

Srivastava, Aarohi; Rastogi, Abhinav; Rao, Abhishek; Shoeb, Abu Awal Md; Abid, Abubakar; Fisch, Adam; Brown, Adam R.; Santoro, Adam; Gupta, Aditya; Garriga-Alonso, Adrià *et al.* 2023. *Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. Transactions on Machine Learning Research*.

Stabler, Edward 1991. Avoid the pedestrian's paradox. In Berwick, Robert C.; Abney, Steven P. & Tenny, Carol (eds.), *Principle-based Parsing: Computation and Psycholinguistics*. Dordrecht: Kluwer. 199-238. <doi.org/10.1007/978-94-011-3474-3_8>.

Stabler, Edward 1997. Derivational minimalism. In Retoré, Christian (ed.), *Logical Aspects of Computational Linguistics*. Berlin / Heidelberg: Springer. 68-95.

Stabler, Edward 2011. Computational Perspectives on Minimalism. In Boeckx, Cedric (ed.), *The Oxford Handbook of Linguistic Minimalism*. Oxford, UK: Oxford University Press. <doi.org/10.1093/oxfordhb/9780199549368.013.0027>.

Stabler, Edward 2013. Two Models of Minimalist, Incremental Syntactic Analysis. *Topics in Cognitive Science* 5,3. 611-633. <doi.org/10.1111/tops.12031>.

Starke, Michal 2001. *Move Dissolves into Merge: A Theory of Locality*. PhD dissertation. Université de Genève.

Steedman, Mark & Baldridge, Jason 2006. Combinatory categorial grammar. In Brown, Keith (ed.), *Encyclopedia of Language & Linguistics*. 2nd edition. Oxford: Elsevier. 610-621.

Steuer, Julius; Mosbach, Marius & Klakow, Dietrich 2023. Large GPT-like Models are Bad Babies: A Closer Look at the Relationship between Linguistic Competence and Psycholinguistic Measures. *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*. Singapore: Association for Computational Linguistics. 114-129. <doi.org/10.18653/v1/2023.conll-babylm.12>.

Stowe, Laurie A.; Kaan, Edith; Sabourin, Laura & Taylor, Ryan C. 2018. The sentence wrap-up dogma. *Cognition* 176. 232-247. <doi.org/10.1016/j.cognition.2018.03.011>.

Strubell, Emma; Ganesh, Ananya & McCallum, Andrew 2019. Energy and policy considerations for deep learning in NLP. In Korhonen, Anna; Traum, David & Màrquez, Lluís (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics. 3645-3650. <DOI: 10.18653/v1/P19-1355>.

Sulger, Sebastian; Butt, Miriam; Holloway King, Tracy; Meurer, Paul; Laczkó, Tibor; Rákosi, György; Bamba Dione, Cheikh M.; Dyvik, Helge; Rosén, Victoria; De Smedt, Koenraad; Patejuk, Agnieszka; Çetinoglu, Özlem; Arka, I Wayan & Mistica, Meladel 2013. ParGramBank: The

ParGram parallel treebank. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, vol. 1. Sofia: Association for Computational Linguistics. 550-560. <www.aclweb.org/anthology/P13-1054.pdf>.

Sutton, Rich 2019. The bitter lesson [Blog post]. *Incomplete Ideas*. <www.incompleteideas.net/IncIdeas/BitterLesson.html>. Last accessed 24/02/2025.

Svenonius, Peter 2016. Significant mid-level results of generative linguistics. <blogg.uit.no/psv000/2016/08/30/significant-mid-level-results-of-generative-linguistics>.

Swanson, Logan 2024. Syntactic learning over tree tiers. In *Proceedings of ESSLLI 2024*. 187-196.

Taylor, Wilson L. 1953. “Cloze Procedure”: A New Tool for Measuring Readability. *Journalism Quarterly* 30,4. 415-433. <doi.org/10.1177/107769905303000401>.

Torr, John 2017. Autobank: A semi-automatic annotation tool for developing deep Minimalist grammar treebanks. In *Proceedings of the demonstrations at the 15th conference of the European chapter of the Association for Computational Linguistics*. 81-86.

Torr, John 2018. Constraining MGbank: Agreement, L-selection and supertagging in minimalist grammars. In Gurevych, Iryna & Miyao, Yusuke (eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. Vol. 1. Melbourne, Australia: Association for Computational Linguistics. 590-600. <DOI: 10.18653/v1/P18-1055>.

Torres, Charles & Futrell, Richard 2023. Simpler neural networks prefer sub-regular languages. In *Findings of the association for computational linguistics: EMNLP 2023*. 1651-1661.

Torres, Charles; Hanson, Kenneth; Graf, Thomas & Mayer, Connor 2023. Modeling island effects with probabilistic tier-based strictly local grammars over trees. In *Proceedings of the Society for Computation in Linguistics (SCI) 2023*. 155-164. <doi.org/10.7275/nz4q-6b09>.

Tran, Tu-Anh & Miyao, Yusuke 2022. Development of a multilingual CCG treebank via Universal Dependencies conversion. In Calzolari, Nicoletta; Béchet, Frédéric; Blache, Philippe; Choukri, Khalid; Cieri, Christopher; Declerck, Thierry; Goggi, Sara; Isahara, Hitoshi; Maegaard, Bente; Mariani, Joseph et al. (eds.), *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association. 5220-5233.

Trinh, Trieu H. & Le, Quoc V. 2019. A simple method for commonsense reasoning. <[arXiv:1806.02847](https://arxiv.org/abs/1806.02847)>.

Trotta, Daniela; Guarasci, Raffaele; Leonardelli, Elisa & Tonelli, Sara 2021. Monolingual and Cross-Lingual Acceptability Judgments with the Italian CoLA corpus. *Findings of the Association for Computational Linguistics: EMNLP 2021*. Punta Cana, Dominican Republic: Association for Computational Linguistics. 2929-2940. <doi.org/10.18653/v1/2021-fnlp-024>.

v1/2021.findings-emnlp.250>.

Turing, Alan M. 1937. Computability and λ -definability. *Journal of Symbolic Logic* 2. 153-163. <doi.org/10.2307/2268280>.

Turing, Alan 1950. Computing machinery and intelligence. *Mind* 59. 433-460. <DOI: 10.1093/mind/lix.236.433>.

van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Oxford University Press. 97-157.

van Riemsdijk, Henk & Williams, Edwin 1986. *Introduction to the Theory of Grammar*. Cambridge, MA: MIT Press.

van Rooij, Iris; Guest, Olivia; Adolfi, Federico; de Haan, Ronald; Kolokova, Antonina & Rich, Patricia 2024. Reclaiming AI as a theoretical tool for cognitive science. *Computational Brain and Behaviour*.

Vaswani, Ashish; Shazeer, Noam; Parmar, Niki; Uszkoreit, Jakob; Jones, Llion; Gomez, Aidan N.; Kaiser, Lukasz & Polosukhin, Illia 2017. Attention Is All You Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S. & Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30 (NIPS 2017)*. Long Beach, CA: Curran Associates, Inc. 5998-6008. <arxiv.org/abs/1706.03762>.

Vermeerbergen, Myriam; Leeson, Lorraine & Crasborn, Onno Alex (eds.) 2007. *Simultaneity in signed languages: Form and function*. Amsterdam: John Benjamins.

Voldoire, A.; Sanchez-Gomez, E.; Salas y Mélia, D.; Decharme, B.; Cassou, C.; Sénési, S.; Valcke, S.; Beau, I.; Alias, A.; Chevallier, M. et al. 2013. The CNRM-CM5.1 global climate model: Description and basic evaluation. *Climate Dynamics* 40.9. 2091-2121. <DOI: 10.1007/s00382-011-1259-y>.

von Humboldt, Wilhelm 1836. *Über die Verschiedenheit des menschlichen Sprachbaues und ihren Einfluß auf die geistige Entwicklung des Menschengeschlechts*. Berlin, Prussia: Druckerei der Königlichen Akademie der Wissenschaften.

Wadler, Philip 1990. Deforestation: Transforming programs to eliminate trees. *Theoretical Computer Science* 73. 231-248. <doi.org/10.1016/0304-3975(90)90147-A>.

Warstadt, Alex & Bowman, Samuel R. 2022. What artificial neural networks can tell us about human language acquisition. In Lappin, Shalom & Bernardy, Jean-Phillipe (eds.), *Algebraic Structures in Natural Language*. Boca Raton: CRC Press, Taylor & Francis. 17-60.

Warstadt, Alex; Mueller, Aaron; Choshen, Leshem; Wilcox, Ethan; Zhuang, Chengxu; Ciro, Juan; Mosquera, Rafael; Paranjape, B.; Williams, A.; Linzen, T. & Cotterell, R. 2023. Findings of the BabyLM Challenge: Sample-Efficient Pretraining on Developmentally Plausible Corpora. *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*. Singapore: Association for Computational Linguistics. 1-6. <doi.org/10.18653/v1/2023.conll-babylm.1>.

Warstadt, Alex; Parrish, Alicia; Liu, Haokun; Mohananey, Anhad; Peng, Wei; Wang, Sheng-Fu & Bowman, Samuel R. 2020. BLiMP: The

Benchmark of Linguistic Minimal Pairs for English. *Transactions of the Association for Computational Linguistics* 8. 377-392. <doi.org/10.1162/tacl_a_00321>.

Warstadt, Alex; Singh, Amanpreet & Bowman, Samuel R. 2018. Neural Network Acceptability Judgments. <[arXiv:1805.12471](https://arxiv.org/abs/1805.12471)>.

Warstadt, Alex; Singh, Amanpreet & Bowman, Samuel R. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics* 7. 625-641. <aclanthology.org/Q19-1040>.

Warstadt, Alex; Zhang, Yian; Li, Xiaocheng; Liu, Haokun & Bowman, Samuel R. 2020. Learning Which Features Matter: RoBERTa Acquires a Preference for Linguistic Generalizations (Eventually). *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 217-235. <doi.org/10.18653/v1/2020.emnlp-main.16>. <aclanthology.org/2020.emnlp-main.16>.

Waskan, Jonathan; Harmon, Ian; Horne, Zachary; Spino, Joseph & Clevenger, John 2014. Explanatory anti-psychologism overturned by lay and scientific case classifications. *Synthese* 191,5. 1013-1035. <DOI: 10.1007/s11229-013-0304-2>.

Wei, Jason; Bosma, Maarten; Zhao, Vincent; Guu, Kelvin; Yu, Adams Wei; Lester, Brian; Du, Nan; Dai, Andrew M. & Le, Quoc V. 2022a. Finetuned language models are zero-shot learners. In *ICLR 2022 Conference Track*. Online: OpenReview

Wei, Jason; Wang, Xuezhi; Schuurmans, Dale; Bosma, Maarten; Ichter, Brian; Xia, Fei; Chi, Ed; Le, Quoc V. & Zhou, Denny 2022b. Chain-of-thought prompting elicits reasoning in large language models. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K. & Oh, A. (eds.), *Advances in Neural Information Processing Systems 35 (NeurIPS 2022) Main Conference Track*. New Orleans, LA: Curran Associates, Inc. 24824-24837.

Wexler, Kenneth & Culicover, Peter W. 1980. *Formal Principles of Language Acquisition*. Cambridge, MA: MIT Press.

Wickelgren, Wayne A. 1969. Context-Sensitive Coding in Speech Recognition, Articulation and Developments. In *Information Processing in The Nervous System: Proceedings of a Symposium Held at the State University of New York at Buffalo 21st-24th October, 1968*. Springer. 85-96.

Wilcox, Ethan; Futrell, Richard & Levy, Roger 2024. Using Computational Models to Test Syntactic Learnability. *Linguistic Inquiry*. 55,4. 805-848. <doi.org/10.1162/ling_a_00491>.

Wilcox, Ethan; Gauthier, Jon; Hu, Jennifer; Qian, Peng & Levy, Roger 2020. On the predictive power of neural language models for human real-time comprehension behavior. In *Proceedings of the Annual Meeting of the Cognitive Science Society*. Online: eScholarship.

Wilcox, Ethan; Levy, Roger; Morita, Takashi & Futrell, Richard 2018. What do RNN Language Models Learn about Filler-Gap Dependencies? In

Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP. Brussels: ACL. 211-221. <arxiv.org/abs/1809.00042>.

Wilkenfeld, Daniel A. 2014. Functional explaining: A new approach to the philosophy of explanation. *Synthese* 191,14. 3367-3391. <DOI: 10.1007/s11229-014-0452-z>.

Wilkenfeld, Daniel A. & Lombrozo, Tania 2020. Explanation classification depends on understanding: Extending the epistemic side-effect effect. *Synthese* 197,6. 2565-2592.

Wilkinson, Mark D.; Dumontier, Michel; Aalbersberg, IJsbrand Jan; Appleton, Gabrielle; Axton, Myles; Baak, Arie; Blomberg, Niklas *et al.* 2016. The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Scientific Data* 3,1. 160018. <doi.org/10.1038/sdata.2016.18>.

Williams, Edwin S. 1977. Discourse and Logical Form. *Linguistic Inquiry* 8,1. 101-139.

Wiltschko, Martina 2008. The syntax of non-inflectional plural marking. *Natural Language and Linguistic Theory* 26,3. 639-694.

Wiltschko, Martina 2014. *The universal structure of categories. Towards a formal typology.* Cambridge: Cambridge University Press.

Wiltschko, Martina 2018. Discovering syntactic variation. In Hornstein, N.; Lasnik, H.; Patel-Grosz, P. & Yang, Ch. (eds.), *Syntactic Structures after 60 Years. The Impact of the Chomskyan Revolution in Linguistics. Studies in Generative Grammar [SGG]* 129. 427-460.

Wiltschko, Martina 2021a. *The grammar of interactional language.* Cambridge: Cambridge University Press.

Wiltschko, Martina 2021b. Universal underpinnings of language-specific categories. A useful heuristic for discovering and comparing categories of grammar and beyond. In Alfieri, Luca; Ramat, Paolo & Arcodia, Giorgio Francesco (eds.), *Linguistic Categories, Language Description and Linguistic Typology*. 59-99.

Wiltschko, Martina 2022. Language is for thought and communication. *Glossa: A Journal of General Linguistics* 7,1. <doi.org/10.16995/glossa.5786>.

Wiltschko, Martina & Heim, Johannes 2016. The syntax of confirmationals. A neo-performative analysis. In Kaltenböck, Gunther; Keizer, Evelien & Lohmann, Arne (eds.), *Outside the Clause. Form and function of extra-clausal constituent.* John Benjamins. 303-340.

Wiltschko, Martina & Heim, Johannes 2020. Grounding Beliefs: Structured Variation in Canadian English Discourse Particles. In Achiri-Taboh, B. (ed.), *Exoticism in English tag questions: Strengthening arguments and caressing the social wheel.* Cambridge: Cambridge Scholars Publishing.

Yang, Andy; Chiang, David & Angluin, Dana 2024. Masked hard-attention transformers recognize exactly the star-free languages. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. &

Zhang, C. (eds.), *Advances in Neural Information Processing Systems 37 (NeurIPS 2024) Main Conference Track*. Vancouver, BC, Canada: Curran Associates, Inc. 10202-10235.

Yang, Charles D. 2016. *The price of linguistic productivity: How children learn to break the rules of language*. Cambridge, MA: MIT Press.

Yang, Yuan & Piantadosi, Steven T. 2022. One model for the learning of language. *Proceedings of the National Academy of Sciences* 119,5. e2021865119. <doi.org/10.1073/pnas.2021865119>.

Yi, Sanghyun; Goel, Rahul; Khatri, Chandra; Cervone, Alessandra; Chung, Tagyoung; Hedayatnia, Behnam; Venkatesh, Anu; Gabriel, Raefer & Hakkani-Tur, Dilek 2019. Towards coherent and engaging spoken dialog response generation using automatic conversation evaluators. In van Deemter, Kees; Lin, Chenghua & Takamura, Hiroya (eds.), *Proceedings of the 12th International Conference on Natural Language Generation*. Tokyo, Japan: Association for Computational Linguistics. 65-75. <DOI: 10.18653/v1/W19-8608>.

Zhang, Chiyuan; Bengio, Samy; Hardt, Mortiz; Recht, Benjamin & Vinyals, Oriol 2021. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM* 64. 107-115. <doi.org/10.1145/3446776>.

Zhang, Yian; Warstadt, Alex; Li, Haau-Sing & Bowman, Samuel R. 2021. When Do You Need Billions of Words of Pretraining Data? In Zong, Chengqing; Xia, Fei; Li, Wenjie & Navigli, Roberto (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing. Vol. 1*. Online: Association for Computational Linguistics. 1112-1125. <DOI: 10.18653/v1/2021.acl-long.90>. <arxiv.org/abs/2011.04946> (2020).

Zhao, M.; Golaz, J. C.; Held, I. M.; Guo, H.; Balaji, V.; Benson, R.; Chen, J. H.; Chen, X.; Donner, L. J.; Dunne, J. P. et al. 2018a. The GFDL global atmosphere and land model AM4.0/LM4.0: 1. Simulation characteristics with prescribed SSTs. *Journal of Advances in Modeling Earth Systems* 10,3. 691-734. <DOI: 10.1002/2017MS001208>.

Zhao, M.; Golaz, J. C.; Held, I. M.; Guo, H.; Balaji, V.; Benson, R.; Chen, J. H.; Chen, X.; Donner, L. J.; Dunne, J. P. et al. 2018b. The GFDL global atmosphere and land model AM4.0/LM4.0: 2. Model description, sensitivity studies, and tuning strategies. *Journal of Advances in Modeling Earth Systems* 10,3. 735-769. <DOI: 10.1002/2017MS001209>.

Zymla, Mark-Matthias 2024. Ambiguity management in computational Glue semantics. In Butt, Miriam; Findlay, Jamie & Toivonen, Ida (eds.), *Proceedings of the LFG'24 Conference*. Konstanz: PubliKon. 285-310. <lfg-proceedings.org/lfg/index.php/main/article/view/59>.

Printed in June 2025
by Industrie Grafiche Pacini Editore Srl
Via A. Gherardesca • 56121 Ospedaletto • Pisa • Italy
Tel. +39 050 313011 • Fax +39 050 3130300
www.pacinieditore.it

