

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A computational operationalisation of competing maturational theories of syntactic development via statistical grammar induction

Permalink

<https://escholarship.org/uc/item/1g35z257>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Marcheva, Mila

Salhan, Suchir

Sun, Weiwei

Publication Date

2026

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

A computational operationalisation of competing maturational theories of syntactic development via statistical grammar induction

Mila Marcheva-Nash

Suchir Salhan

Weiwei Sun

mmm67@cam.ac.uk

sas2450@cam.ac.uk

ws390@cam.ac.uk

Department of Computer Science & Technology, University of Cambridge

Abstract

This paper is concerned with what intermediate syntactic categories children acquire during first language development, and in what order. Maturational theories make different predictions. Bottom-up accounts (GROWING) propose that lexical and inflectional structure emerges first, while inward accounts (INWARD) predict early access to discourse-related categories. We computationally operationalise these hypotheses of staged syntactic emergence using statistical grammar induction, asking what each proposed ordering makes learnable when input and learning algorithm are held constant. Our framework makes category acquisition explicit and allows us to explore how different maturational orderings shape the structure that can be learned under identical conditions. Based on this operationalisation, the GROWING account significantly outperforms the INWARD account across three evaluation metrics.

Keywords: language acquisition; syntactic development; grammar induction; maturation

Introduction

A central question in first language acquisition (FLA) is how children develop an adult-like grammatical system (Brown, 1973). Linguistic theories differ on whether grammatical categories are predetermined by biology or emerge gradually from experience. Continuity approaches within the generative tradition assume that all categories are innate and available from birth, consistent with the idea of Universal Grammar (UG) (Pinker, 1984; Boser, Lust, Santelmann, & Whitman, 1992; Clahsen, Eisenbeiss, & Vainikka, 1994). Maturational accounts assume that certain syntactic categories are innate but become accessible only at specific points in development, shaping the order of acquisition (Radford, 1988, 1990; Guasti, 1993). By contrast, the emergentist perspective emphasises that categories emerge from patterns in the input under cognitive constraints (O’Grady, 2008; Bosch & Biberauer, 2024). Functionalist, usage-based, and constructivist theories are in line with emergentism (Bates & MacWhinney, 1987; Tomasello, 2005; Lieven & Tomasello, 2008; Ambridge & Lieven, 2015; Lieven, 2016; Behrens, 2021). What is common across maturational accounts and emergentist accounts is the rejection of continuity, i.e. both theories agree that adult grammatical competence is not available from the outset (Tomasello, 2000; Bosch, 2025). What maturation and emergentism disagree on is the source for the staged development: under maturation, the order in which the categories appear is innately encoded, whereas under emergentism, the

order is determined by the interplay of input and cognitive constraints.

Within the maturational tradition theories differ sharply on which categories appear first. Bottom-up proposals, such as the GROWING Trees Hypothesis (Friedmann, Belletti, & Rizzi, 2021), suggest that lexical and inflectional categories (N, V, T) emerge first, followed by discourse-related functional categories. Inward maturation proposals, like the INWARD Growing Spine Hypothesis (Heim & Wiltchko, 2025), predict early access to higher discourse-related categories, with lower categories appearing later. These hypotheses formalise patterns observed in corpora of child and child-directed speech. Computational modelling can complement traditional methods by allowing us to explore whether the proposed developmental orderings make different grammatical structures recoverable from the same input.

In this paper, we introduce a computational framework for testing competing theories of staged syntactic category acquisition under developmentally plausible constraints. In grammar induction, the statistical learner receives strings as input and infers explicit hierarchical structures, a grammar, making the order of category acquisition observable rather than stipulated (Johnson, 2013). By controlling the order in which categories become accessible, we simulate the developmental trajectories predicted by different maturational theories, while holding constant the input and the learning algorithm. Comparing the resulting induced grammars across controlled conditions reveals how different maturational orderings constrain what is learnable from identical input and under identical conditions.

We present an experimental suite to compare two maturational hypotheses, GROWING and INWARD. We use morphemically tokenised child-directed speech as developmentally plausible input, reflecting linguistic units to which children are known to be sensitive (Pearl & Sprouse, 2013; Marcheva, Biberauer, & Sun, 2025). Across a wide range of conditions, we find empirical advantages for staged category acquisition, particularly under the GROWING curriculum, relative to Continuity approaches. While both GROWING and INWARD curricula converge to comparable global performance, GROWING significantly outperforms INWARD on all queried metrics (F1, Jensen-Shannon divergence, and child speech sentence log-likelihood). At the level of the learning process, GROWING favours earlier stabilisation of

phrasal structure, whilst INWARD yields lower divergence for certain clause-level categories. Together, these results align with the qualitative predictions of the corresponding maturational accounts and demonstrate the utility of our method for systematically exploring alternative trajectories of staged syntactic development. The framework and the tested staged grammars are available on GitHub¹. While we pilot the framework on two maturational hypotheses, the framework could be extended to emergentist hypotheses, given appropriate staged curricula are developed.




Figure 1: Pipeline for statistical learner modelling staged syntactic development, explained in detail in **Methodology**.

Background

Syntactic development

A long tradition in FLA characterises syntactic development as staged, following a trajectory of single-word utterance, then unmarked two-word phrases, and ultimately phrases marked with increasingly complex morphosyntactic structure (Brown, 1973). Generative accounts of syntactic acquisition centre on when and how functional categories are acquired. There are two main schools of thought, with regards to when functional categories, assumed to be part of Universal Grammar (UG), become available. **Continuity** (Westergaard, 2009) posits that all information in UG is available from

the start. Thus, the functional structure of children’s initial grammar is not significantly different from adult’s grammars. **Maturation** (Friedmann et al., 2021; Heim & Wilschko, 2025) posits that grammatical categories and principles from UG are not fully available to children initially, but instead are incrementally accessed. Thus, UG specifies not only the hierarchical structure and functional categories, but the order in which they become available to the acquirer, too. Under maturation the UG-given hierarchical structure is fixed, but children gradually gain more access to parts of the structure. All of the above hypotheses assume that functional categories have fixed granularity. Thus, for completion we must also mention neo-emergentism, a hypothesis of syntactic acquisition which posits that increasing (flexible) granularity of the functional categories is a key aspect of syntactic acquisition (Biberauer & Ian, 2015; Bosch & Biberauer, 2024; Bosch, 2025).

With the distinction of order of acquisition, maturation can be further broken down into bottom-up and inward orders (w.r.t. to the UG predefined spine). Bottom-up maturation posits that the categories which become available first are the ones closer to the leaves. The most recent bottom up approach, which also incorporates cartography, is the **GROWING Trees Hypothesis** (Friedmann et al., 2021). It distinguishes between 3 stages: in stage 1 only IP/TP and VP are available (allowing for inflection and A(argument)-movement); in stage 2 the lower left-periphery (e.g. allowing for wh- questions) becomes available; finally in stage 3 the entire cartographic hierarchy becomes available, including topicalisations and embeddings. Inward maturation, of which an example is the **INWARD Growing Spine Hypothesis** (Heim & Wilschko, 2025), postulates the early development of CP, which appears at the middle of the cartographic spine.

Grammar induction

Grammar induction (GI) is the process of learning the hierarchical structure which is latent in language data. A grammar is composed of a **symbolic component** (the rules) and a **probabilistic component** (the probabilities assigned to the rules). A subtype of GI is grammar reestimation (Pereira & Schabes, 1992), where the symbolic component is provided, and the probabilistic component needs to be learned. Foundational work in GI relies on statistical methods (Klein & Manning, 2004; Johnson, Griffiths, & Goldwater, 2007; Carroll & Charniak, 1992; Bisk & Hockenmaier, 2012; Clark & Schuler, 2023). In recent years neural grammar induction has been put forward as a method which induces grammars with unprecedented F1 scores from raw data (Kim, Dyer, & Rush, 2019). However, neural methods are less interpretable than the foundational statistical models. Thus, for this paper we rely on statistical GI, specifically grammar reestimation, with a PCFG.

Using GI as an approximation for language acquisition, and specifically syntax acquisition, is well motivated in the literature (Johnson, 2013); examples include

¹<https://github.com/milamarcheva/maturational-grammar-induction>

Stage	Description	Penn Treebank tags
BASE GROWING	Simple NPs and VPs, no other phrasal material	NTs: ROOT, S, FRAG, NP, VP PTs: NN, NNP, PRP, PRP\$, VB
BASE INWARD	Discourse-level combinations (e.g. INTJ $\rightarrow$ INTJ INTJ)	PTs: NN, PRP, PRP\$, VB, UH NTs: ROOT, S, FRAG, INTJ
VP	Predicate and modifier structure	NTs: ADJP, ADVP, PP, NNPS, NNS PTs: DT, PDT, POS, NNS, CD, JJ, JJR, JJS, RB, RBR, RBS, IN, RP, VBG, VBN, NOT, DIV
TP	Finite/tense and auxiliary marking	PTs/NTs: VBD, VBP, VBZ, AUX, COP, MD, ASP, T, PRS
CP	Complementisers and operators	PTs: COMP, CC, WP, WP\$, WRB, WDT NTs: SBAR, SBARQ, SINV, SQ, WHADJP, WHADVP, WHNP
INTJ	Interjections	PTs: UH NTs: INTJ

Table 1: Syntactic stages are defined as cumulative sets of Penn Treebank categories. Note that in the original PTB some tags are PTs, but here are NTs due to the morphemic tokenisation (e.g. NNS, VBG).

i.a. Buttery (2004); Perfors, Tenenbaum, and Regier (2011); Kwiatkowski, Goldwater, Zettlemoyer, and Steedman (2012)). A key advantage of GI over traditional linguistic analyses, which often target isolated phenomena, is that it provides a unified account of all sentences in a corpus within a single model. Note that GI remains an exploratory tool and does not capture the full complexity of first language acquisition, as it abstracts away from non-language cues (Dupre, 2024).

Probabilistic Context-Free Grammar (PCFG) A PCFG is defined as $G = (NT, \Sigma, R, S, F)$, where NT is the set of non-terminals, Σ the vocabulary, R the rules, S the start symbol, and F the rule-probability function. We distinguish production rules from lexicalisations, since the former are the main object of staged syntactic access.

Variational Bayes (VB) Statistical GI standardly uses one of two estimation methods: Expectation Maximisation (EM) (Klein & Manning, 2002) or Variational Bayes (VB) (Kurihara & Sato, 2006). EM is a frequentist procedure that estimates rule probabilities solely from expected counts in the data, whereas VB is a Bayesian extension of EM that introduces prior distributions over rule probabilities, most commonly Dirichlet priors (Johnson et al., 2007; Liang, Petrov, Jordan, & Klein, 2007). In VB, each grammar rule is associated with a pseudocount parameter (α) that influences learning at every iteration, encoding prior confidence in a rule. Higher values bias the model toward retaining a rule, while lower values allow unsupported rules to shrink in probability. Pseudocounts allow probabilities learned in earlier stages to be carried forward as priors, while still permitting new rules to be learned.

Methodology

Figure 1 illustrates our computational operationalisation of staged syntactic development: maturational hypotheses are translated into curricula comprised of rule subsets of a symbolic grammar, and acquisition is approximated as Bayesian grammar induction.

Acquisition-inspired curricula

In order to translate the theoretical claims about maturation and the timing of CP into a curriculum-based grammar induction problem, we define two curricula approximating the incremental access to syntactic categories stipulated by GROWING and INWARD maturation hypotheses. Previously Salhan, Diehl Martinez, Goriely, and Buttery (2024) have approached this problem using universal part-of-speech (UPOS) tags (Petrov, Das, & McDonald, 2012), however, we instead rely on the Penn Treebank (PTB) tagset (Marcus, Mitchell P., Santorini, Beatrice, & Mary Ann Marcinkiewicz, 1995), which is more detailed. The PTB tagset encodes both phrase type (e.g. VP, NP, PP, INTJ) and the presence of functional material (e.g. AUX, MD, TO, COP, complementisers, wh-phrases), which allows us to define more fine-grained cognitively-inspired curricula. Stages are defined cumulatively over Penn Treebank tags in Table 1. Note that the stages we define are an approximation of the maturational theories, and alternative definitions of the stages can be substituted within the same pipeline.

The two maturational curricula, constructed by ordering the stages from Table 1 are listed below, as well as an explicit statement of the CONTINUITY condition:

- **GROWING:** BASEGROWING, VP, TP, CP, INTJ
- **INWARD:** BASEINWARD, BASEGROWING, CP, TP, VP
- **CONTINUITY:** all rules are available from the start

Maturational syntactic development via VB

Maturational syntactic development poses that syntactic knowledge becomes accessible incrementally in distinct stages (refer to Table 1 for a PTB approximation of these stages). Using VB we can perform learning in stages, where increasingly more rules become available. As new rules are being introduced at every stage, pseudocounts allow probabilities learned in earlier stages to be carried forward as priors, while still permitting new rules to be learned.

We formulate the pseudocount of existing and new rules in Eq. 1, which allows information to be carried across stages as the learning space expands. After the completion of stage k , the posterior mean rule probabilities $\mathbb{P}^k$ are converted into a

Dirichlet prior vector α^{k+1} for production rules $x \in X$ in stage $k+1$ according to:

$$\alpha_{i,x}^{k+1} = \begin{cases} N^k s_p p_i^k + 0.1, & 1 \leq i \leq N^k \\ \frac{N^k s_p \eta}{X^{k+1} - X^k} + 0.1, & N^k < i \leq N^{k+1} \end{cases} \quad (1)$$

where N^k is the number of sentences successfully parsed at stage k , s_p is a scaling parameter for production rule priors carried forward from the previous stage, and η is a dial for how much probabilistic mass can be allotted to newly introduced rules in a stage. X^k is the number of rules which share the same NT on the LHS as x in stage k , so the mass for new rules is distributed across the newly available expansions for that NT. Lexicalisation priors are treated analogously, using a separate scaling parameter s_ℓ .² Thus, s_p and s_ℓ control the degree of inter-stage memory, while η controls how readily the learner assigns probability mass to newly available rules. When $s = 0$, learning from the previous stage is ignored, whereas larger values increasingly bias the learner toward retaining previously learned distributions.

Experimental setup

System

We use the ‘‘Inside-Outside algorithm for estimating PCFGs from terminal strings’’ (Johnson, 2007), which iteratively estimates rule probabilities by optimising likelihood-based objectives from unlabelled strings input. Our setup is based on Mark Johnson’s implementation, available on his website³. This is an example of a grammar reestimation system, which requires as input the symbolic component of a PCFG, i.e. the rules, initialised with a weight and a pseudocount, as well as a list of sentences based on which the probability distributions of the available rules are induced. We manipulate the pseudocount as discussed in the **Methodology** section. We select to run experiments with 20 iterations because hyperparameter tuning shows plateauing around this point.

Data

We use the child-directed speech sentences and parses from the morphemically tokenised version of CHILDES-TB (Pearl & Sprouse, 2013; Marcheva et al., 2025), which we restrict to *sentence length* > 1. The dataset amounts to 126,152 sentences and their corresponding parses. CHILDES-TB⁴ is based on the child-directed speech from five corpora, covering 110 children aged from 6 months to 6 years.

ORACLE grammar

We extract a PCFG from the morphemically tokenised CHILDES-TB parses which we use as an ORACLE (gold-

standard) grammar in evaluation, whilst its symbolic component serves as the initial grammar for the GI system. The ORACLE grammar is a PTB-style PCFG approximating adult competence. We do not claim that this PCFG is an absolute representation of adult mental grammar; rather, it provides a linguistically interpretable target against which different staged access hypotheses can be compared. We restrict production rules by minimum frequency and select the smallest grammar that preserves 100% parse coverage. At $f_m = 7$, the ORACLE grammar is comprised of 1,387 production rules and 8,563 lexicalisations over a vocabulary of 6,273 words.

Evaluation

In addition to using the CHILDES-TB parses as a gold standard, based on which we calculate unlabelled F1 score, we also present distributional metrics which allow for the comparison of the distributions of the learned grammars to the distribution of the ORACLE grammar.

We use an induced grammar to parse 1,000 randomly selected parses from CHILDES-TB. Then we calculate **unlabelled F1**, implemented in line with PARSEVAL (Black et al., 1991), on parses from the induced grammar and the gold parses from CHILDES-TB.

We evaluate the induced grammars as generative language models using mean length-normalised **marginal log-likelihood** (see Eq. 2) over 91,901 child speech sentences from the corpora used in CHILDES-TB, restricted to the lexicon of the ORACLE grammar.⁵ Scores are computed only for utterances licensed by at least one complete parse; larger values are better. For a PCFG G , let $\mathcal{T}(\mathbf{w}^{(i)})$ be the set of all parse trees for an utterance $\mathbf{w}^{(i)}$:

$$\overline{\log p_{\text{sent}}} = \frac{1}{N} \sum_{i=1}^N \frac{\ell_i}{T_i}, \quad \ell_i = \ln \sum_{t \in \mathcal{T}(\mathbf{w}^{(i)})} P_G(t) \quad (2)$$

We also compare the rule expansion distributions of ORACLE and induced grammars for each non-terminal A using **Jensen-Shannon divergence (JSD)** (Lin, 1991), as formulated in Eq. 3. A result of $JSD = 0$ indicates identical distributions.

$$JSD(p_A, q_A) = \frac{1}{2} \text{KL}(p_A \| m_A) + \frac{1}{2} \text{KL}(q_A \| m_A), \quad (3)$$

where $m_A = \frac{1}{2}(p_A + q_A)$.

Results

Continuity vs. Maturation

Under the CONTINUITY condition, all grammar rules are available from the outset. In contrast, under MATURATION, rules are introduced in discrete stages, and the probabilities learned at each stage are carried forward to the next via pseudocount transfer, as defined in Eq. 1. Table 2 provides an

²Treating lexicalisations and productions with separate pseudocount values allows us to focus specifically on the production rules as they are more pertinent to the topic of syntactic development.

³<https://web.science.mq.edu.au/~mjohnson/Software.htm>

⁴<https://sites.socsci.uci.edu/~lpearl/CoLaLab/CHILDESTreebank/>

⁵A larger child-speech test set could be constructed without this lexical restriction. The age range and number of children match CHILDES-TB.




Figure 2: GROWING vs INWARD learning progression across the 5 learning stages, compared via F1, $\overline{\log p_{\text{sent}}}$, and JSD. GROWING consistently outperforms INWARD. For reference, the ORACLE grammar baseline for $\overline{\log p_{\text{sent}}} = -6.1823$ is provided. Higher F1 and $\overline{\log p_{\text{sent}}}$, and lower JSD, indicate better performance.

overview of the F1 scores achieved by systems approximating the three syntactic development hypotheses we operationalised: CONTINUITY, GROWING, INWARD.

CONTINUITY experiments are run for 20 Variational Bayes iterations under a Dirichlet prior with production and lexicalisation pseudocounts $\alpha_p, \alpha_\ell \in \{0.1, 0.3, 0.5\}$. Across these settings, performance is stable: mean F1 (see also Table 2) is 0.799 ± 0.001 . We adopt $\alpha_p = \alpha_\ell = 0.1$ as the CONTINUITY baseline for comparison with the MATURATION conditions, reflecting our choice of minimal constant pseudocounts (Eq. 1). Under this setting, performance is $F1 = 0.8000$, $\overline{\log p_{\text{sent}}} = -6.0285$, and $JSD = 0.1928$. With appropriate hyperparameter settings, the GROWING curriculum outperforms the CONTINUITY condition (see Table 2), whereas the INWARD curriculum consistently underperforms.

Condition	F1
CONTINUITY $\alpha_p, \alpha_\ell \in \{0.1, 0.3, 0.5\}$ mean, st.d.	0.7994, .0013
CONTINUITY ($\alpha_p = \alpha_\ell = 0.1$)	0.8000
GROWING mean, st.d.	0.8059, .0038
INWARD mean, st.d.	0.7953, .0022
Neural GI	0.7856

Table 2: F1 scores comparing CONTINUITY and MATURATION. Neural GI baseline from Table 2 in Marcheva et al. (2025) is provided for comparison.

GROWING vs. INWARD

Final grammar comparison We compare final grammars, after stage 5, induced under the GROWING and INWARD curricula. We report the top 10 hyperparameter settings (ranked by GROWING F1) in Table 4; in all cases, GROWING also exceeds the CONTINUITY baseline ($\alpha_p = \alpha_\ell = 0.1$, $F1 = 0.8000$). Paired, one-sided Wilcoxon signed-rank tests over 72 matched hyperparameter configurations spanning $s_\ell \in \{0.001, 0.01, 0.1\}$, $s_p \in \{0.001, 0.005, 0.01, 0.05, 0.1, 0.2\}$, and $\eta \in \{0.001, 0.005, 0.01, 0.05\}$ show that GROWING significantly outperforms INWARD on all evaluation metrics (see Table 3). The underperformance of INWARD may be explained either by the early introduction of INTJ, or by the

very late introduction of VP. As observed in Figure 2, the final stage for INWARD, corresponding to the VP stage in Table 1, leads to the largest improvement in performance across all three metrics. This makes it more likely that the underperformance of INWARD is due to the late introduction of VP. This implies that GROWING better captures utterances requiring predicate, argument, and modifier structure, whereas any INWARD advantage is restricted to interactional or clause-level material.

Metric	Median Δ	p -value	$r_{\text{rank-biserial}}$	Direction
F1	+0.0056	1.5×10^{-12}	0.95	GROWING
$\overline{\log p_{\text{sent}}}$	+0.0104	7.4×10^{-13}	0.96	GROWING
JSD	-0.0198	8.3×10^{-14}	-1.00	GROWING

Table 3: Results from a paired one-sided Wilcoxon signed-rank test comparing 72 pairs of final grammars induced under GROWING and INWARD.

Learning progression The learning progression through the stages of the two maturational curricula is illustrated in Figure 2. The F1 under GROWING is consistently higher throughout the stages. Furthermore, the learning progression of the GROWING curriculum for all metrics indicates monotonic improvement, in contrast with INWARD where some stages decrease performance (e.g. stage 1-2 for $\overline{\log p_{\text{sent}}}$ or stage 2-3 for mean JSD). For INWARD’s F1 score, the last stage is primarily responsible for the final F1 achieved. This corresponds to the VP stage, which introduces argument movement. The mean JSD reflects a similar pattern to F1. The plot of the mean normalised sentence log-likelihood, $\overline{\log p_{\text{sent}}}$, shows that both maturational conditions lead to improvement over the oracle (initial grammar). This result affirms that the maturational hypotheses produce grammars more favourable to explaining child productions.

JSD per phrase NT We examine the stage-wise progression of JSD for selected phrase-level NTs introduced at different points under the GROWING and INWARD curricula. Tracking individual categories allows us to compare how




Figure 3: Mean JSD for instrumental phrase-level NTs. GROWING is in blue and INWARD is in orange, but the different NTs are illustrated with different figure markers and line styles. Low scores of JSD indicate similarity with the ORACLE distribution.

s_p	η	GROWING			INWARD		
		F1	$\log p_{\text{sent}}$	JSD	F1	$\log p_{\text{sent}}$	JSD
0.01	0.001	0.8121	-6.0545	0.2558	0.7935	-6.0941	0.2739
0.01	0.05	0.8106	-6.0542	0.2548	0.7943	-6.0941	0.2727
0.01	0.01	0.8103	-6.0543	0.2555	0.7943	-6.0941	0.2736
0.01	0.005	0.8102	-6.0544	0.2557	0.7936	-6.0941	0.2738
0.005	0.005	0.8080	-6.0547	0.2572	0.7927	-6.0929	0.2730
0.005	0.01	0.8078	-6.0546	0.2572	0.7929	-6.0929	0.2730
0.001	0.001	0.8077	-6.0543	0.2584	0.7976	-6.0908	0.2726
0.001	0.01	0.8070	-6.0543	0.2583	0.7978	-6.0908	0.2726
0.005	0.001	0.8070	-6.0547	0.2573	0.7929	-6.0929	0.2730
0.001	0.005	0.8068	-6.0543	0.2584	0.7974	-6.0908	0.2726

Table 4: Top-performing hyperparameter setups for **GROWING** and **INWARD** curricula. $s_\ell = 0.1$ for all of these. Higher F1 and $\log p_{\text{sent}}$, and lower JSD, indicate better performance.

alternative maturational orderings affect the stabilisation of specific components of the grammar.

Under GROWING, core phrasal categories such as NP and VP show gradual, largely monotonic reductions in JSD, with a similar pattern for VP-stage modifier phrases (PP, ADVP, ADJP). Under INWARD, these VP-related categories, which are introduced only at the final stage, exhibit higher JSD relative to GROWING, suggesting as expected, that the more stages where the category is active would lead to a more refined result. For clause-level categories, including INTJ and WH-related phrases, the two curricula show more mixed trajectories, with no uniform advantage for either ordering. By making the emergence of NTs explicit, we provide a quantitative tool for exploring how different developmental orderings shape learning dynamics.

Discussion

While our experiments showcase a flexible computational framework for comparing syntactic development theories under controlled conditions, several limitations of the present study should be acknowledged. First, our evaluation relies on CHILDES-TB, which provides manually annotated parses for English CDS. Although the GI framework itself is not language-specific, extending this approach to other languages requires the appropriately annotated resources. Second, the curricula may differ in the number of rules available at each

developmental stage. Future work could further control for this factor by testing more fine-grained curricula. Third, the type of GI our framework relies on is grammar reestimation. The symbolic part of the grammar needs to be provided, and learning consists of estimating the probabilistic component of the grammar from the input, which consists of unlabelled sentences. This design choice ensures a controlled comparison between maturational hypotheses but limits the model’s ability to discover novel structures.

Although we operationalise two maturational hypotheses, GROWING and INWARD, our pipeline for staged syntactic development, illustrated in Figure 1 is not tied to maturation as the source for ordering in syntactic development. Usage-based and constructivist theories also reject continuity, i.e. that children have adult grammatical knowledge from the outset (Tomasello, 2000, 2005; Behrens, 2021). The main difference between maturational and usage-based accounts is the source of the order of syntactic development: under maturation, the order of categories is innate, whereas under usage-based accounts it emerges constrained by the input and domain-general cognitive capacities of the child. Our pipeline requires that the stages of syntactic development are pre-specified before training. Thus, it can be used for non-generativist accounts, if a curriculum is specified based on item-specific frames or constructions. In that way, our staged grammar-induction paradigm provides a general method for comparing theories that reject the CONTINUITY hypothesis, even if they disagree about why the learner’s hypothesis space changes over time.

Conclusion

We present a grammar induction framework for staged syntactic development hypotheses, which makes the emergence of syntactic categories explicit and enables a controlled comparison of alternative developmental orderings. We operationalise two competing maturational hypotheses of syntactic development, GROWING and INWARD, and find that under our framework GROWING significantly outperforms INWARD. This pilot comparison illustrates how our framework can serve as a principled tool for investigating alternative staged hypotheses of syntactic development.

References

- Ambridge, B., & Lieven, E. (2015, January). *A constructivist account of child language acquisition*. Wiley. Retrieved from <http://dx.doi.org/10.1002/9781118346136.ch22> doi: 10.1002/9781118346136.ch22
- Bates, E., & MacWhinney, B. (1987). Competition, variation, and language learning. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 157–193). Lawrence Erlbaum.
- Behrens, H. (2021, August). Constructivist approaches to first language acquisition. *Journal of Child Language*, 48(5), 959–983. Retrieved from <http://dx.doi.org/10.1017/S0305000921000556> doi: 10.1017/S0305000921000556
- Biberauer, T., & Ian, R. (2015, 01). Rethinking formal hierarchies: A proposed unification. *Beyond Functional Sequence*, 10, 295–313. Retrieved from https://www.mml1.cam.ac.uk/files/copil_7.1_biberauerroberts.pdf
- Bisk, Y., & Hockenmaier, J. (2012, June). Induction of linguistic structure with Combinatory Categorical Grammars. In T. Cohn, P. Blunsom, & J. Graca (Eds.), *Proceedings of the NAACL-HLT workshop on the induction of linguistic structure* (pp. 90–95). Montréal, Canada: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/W12-1912>
- Black, E., Abney, S., Flickinger, D., Gdaniec, C. H., Grishman, R., Harrison, P., ... Strzalkowski, T. (1991). A procedure for quantitatively comparing the syntactic coverage of english grammars. In *Proceedings of the february 1991 darpa speech and natural language workshop* (pp. 306–311). Morgan Kaufmann.
- Bosch, N. (2025). Categorical granularity in syntactic acquisition: A multilingual corpus study on the left periphery. *Glossa*, 10. doi: 10.16995/glossa.23986
- Bosch, N., & Biberauer, T. (2024). Emergent Syntactic Categories and Increasing Granularity: Evidence from a Multilingual Corpus Study. In *Proceedings of bucl48*.
- Boser, K., Lust, B., Santelmann, L., & Whitman, J. (1992). The syntax of cp and v-2 in early child german (ecg): The strong continuity hypothesis. In *Proceedings of the north-east linguistic society (nels)* 22 (pp. 51–66). Amherst: University of Massachusetts.
- Brown, R. (1973). Development of the first language in the human species. *American psychologist*, 28(2), 97.
- Buttery, P. (2004, August 28th). A quantitative evaluation of naturalistic models of language acquisition; the efficiency of the triggering learning algorithm compared to a categorical grammar learner. In W. G. Sakas (Ed.), *Proceedings of the workshop on psycho-computational models of human language acquisition* (pp. 3–10). Geneva, Switzerland: COLING. Retrieved from <https://aclanthology.org/W04-1301/>
- Carroll, G., & Charniak, E. (1992). *Two experiments on learning probabilistic dependency grammars from corpora* (Tech. Rep. No. CS-92-16). Providence, RI: Dept. of Computer Science, Brown University. Retrieved from <ftp://wilma.cs.brown.edu/pub/techreports/92/cs92-16.ps.Z>
- Clahsen, H., Eisenbeiss, S., & Vainikka, A. (1994). The seeds of structure: A syntactic analysis of the acquisition of case marking. In T. Hoekstra & B. Schwartz (Eds.), *Language acquisition studies in generative grammar* (pp. 85–118). Amsterdam: John Benjamins. doi: 10.1075/lald.8.07cla
- Clark, C., & Schuler, W. (2023, July). Categorical grammar induction from raw data. In A. Rogers, J. Boyd-Graber, & N. Okazaki (Eds.), *Findings of the association for computational linguistics: Acl 2023* (pp. 2368–2379). Toronto, Canada: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/2023.findings-acl.149> doi: 10.18653/v1/2023.findings-acl.149
- Dupre, G. (2024, June). Acquiring a language vs. inducing a grammar. *Cognition*, 247, 105771. Retrieved from <http://dx.doi.org/10.1016/j.cognition.2024.105771> doi: 10.1016/j.cognition.2024.105771
- Friedmann, N., Belletti, A., & Rizzi, L. (2021, November). Growing trees: The acquisition of the left periphery. *Glossa: a journal of general linguistics*, 6(1). Retrieved from <http://dx.doi.org/10.16995/glossa.5877> doi: 10.16995/glossa.5877
- Guasti, M. T. (1993). Verb syntax in italian child grammar: Finite and nonfinite verbs. *Language Acquisition*, 3(1), 1–40. doi: 10.1207/s15327817la0301_1
- Heim, J., & Wiltschko, M. (2025, May). Rethinking structural growth: Insights from the acquisition of interactional language. *Glossa: a journal of general linguistics*, 10(1). Retrieved from <http://dx.doi.org/10.16995/glossa.16396> doi: 10.16995/glossa.16396
- Johnson, M. (2007, June). Why doesn't EM find good HMM POS-taggers? In J. Eisner (Ed.), *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)* (pp. 296–305). Prague, Czech Republic: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/D07-1031/>
- Johnson, M. (2013). Language acquisition as statistical inference. *The language-cognition interface*, 109–134.
- Johnson, M., Griffiths, T., & Goldwater, S. (2007, April). Bayesian inference for PCFGs via Markov chain Monte Carlo. In C. Sidner, T. Schultz, M. Stone, & C. Zhai (Eds.), *Human language technologies 2007: The conference of the north American chapter of the association for computational linguistics; proceedings of the main conference* (pp. 139–146). Rochester, New York: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/N07-1018/>
- Kim, Y., Dyer, C., & Rush, A. (2019, July). Compound

- probabilistic context-free grammars for grammar induction. In A. Korhonen, D. Traum, & L. Màrquez (Eds.), *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 2369–2385). Florence, Italy: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/P19-1228> doi: 10.18653/v1/P19-1228
- Klein, D., & Manning, C. (2004, July). Corpus-based induction of syntactic structure: Models of dependency and constituency. In *Proceedings of the 42nd annual meeting of the association for computational linguistics (ACL-04)* (pp. 478–485). Barcelona, Spain. Retrieved from <https://aclanthology.org/P04-1061> doi: 10.3115/1218955.1219016
- Klein, D., & Manning, C. D. (2002, July). A generative constituent-context model for improved grammar induction. In P. Isabelle, E. Charniak, & D. Lin (Eds.), *Proceedings of the 40th annual meeting of the association for computational linguistics* (pp. 128–135). Philadelphia, Pennsylvania, USA: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/P02-1017/> doi: 10.3115/1073083.1073106
- Kurihara, K., & Sato, T. (2006). Variational bayesian grammar induction for natural language. In *Grammatical inference: Algorithms and applications* (p. 84–96). Springer Berlin Heidelberg. Retrieved from http://dx.doi.org/10.1007/11872436_8 doi: 10.1007/11872436_8
- Kwiatkowski, T., Goldwater, S., Zettlemoyer, L., & Steedman, M. (2012, April). A probabilistic model of syntactic and semantic acquisition from child-directed utterances and their meanings. In W. Daelemans (Ed.), *Proceedings of the 13th conference of the European chapter of the association for computational linguistics* (pp. 234–244). Avignon, France: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/E12-1024/>
- Liang, P., Petrov, S., Jordan, M., & Klein, D. (2007, June). The infinite PCFG using hierarchical Dirichlet processes. In J. Eisner (Ed.), *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)* (pp. 688–697). Prague, Czech Republic: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/D07-1072/>
- Lieven, E. (2016). Usage-based approaches to language development: Where do we go from here? *Language and Cognition*, 8(3), 346–368. doi: 10.1017/langcog.2016.16
- Lieven, E., & Tomasello, M. (2008). Children’s first language acquisition from a usage-based perspective: In p.robinson and n.ellis (eds). In *Handbook of cognitive linguistics and second language acquisition* (pp. 168–196).
- Lin, J. (1991). Divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory*, 37(1), 145–151. Retrieved from <http://dx.doi.org/10.1109/18.61115> doi: 10.1109/18.61115
- Marcheva, M., Biberauer, T., & Sun, W. (2025, May). Profiling neural grammar induction on morphemically tokenised child-directed speech. In T. Kuribayashi, G. Rambelli, E. Takmaz, P. Wicke, J. Li, & B.-D. Oh (Eds.), *Proceedings of the workshop on cognitive modeling and computational linguistics* (pp. 47–54). Albuquerque, New Mexico, USA: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/2025.cmcl-1.7/> doi: 10.18653/v1/2025.cmcl-1.7
- Marcus, Mitchell P., Santorini, Beatrice, & Mary Ann Marcinkiewicz. (1995). *Treebank-2*. Linguistic Data Consortium. Retrieved from <https://catalog.ldc.upenn.edu/LDC95T7> doi: 10.35111/WF9P-G717
- O’Grady, W. (2008, April). The emergentist program. *Lingua*, 118(4), 447–464. Retrieved from <http://dx.doi.org/10.1016/j.lingua.2006.12.001> doi: 10.1016/j.lingua.2006.12.001
- Pearl, L. S., & Sprouse, J. (2013). Syntactic islands and learning biases: Combining experimental syntax and computational modeling to investigate the language acquisition problem. *Language Acquisition*. Retrieved from <https://ling.auf.net/lingbuzz/001493>
- Pereira, F., & Schabes, Y. (1992, June). Inside-outside reestimation from partially bracketed corpora. In *30th annual meeting of the association for computational linguistics* (pp. 128–135). Newark, Delaware, USA: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/P92-1017/> doi: 10.3115/981967.981984
- Perfors, A., Tenenbaum, J. B., & Regier, T. (2011, March). The learnability of abstract syntactic principles. *Cognition*, 118(3), 306–338. Retrieved from <https://doi.org/10.1016/j.cognition.2010.11.001> doi: 10.1016/j.cognition.2010.11.001
- Petrov, S., Das, D., & McDonald, R. (2012, May). A universal part-of-speech tagset. In N. Calzolari et al. (Eds.), *Proceedings of the eighth international conference on language resources and evaluation (LREC’12)* (pp. 2089–2096). Istanbul, Turkey: European Language Resources Association (ELRA). Retrieved from <https://aclanthology.org/L12-1115/>
- Pinker, S. (1984). *Language learnability and language development*. Cambridge: Cambridge University Press.
- Radford, A. (1988). *Transformational grammar: A first course*. Cambridge University Press. Retrieved from <http://dx.doi.org/10.1017/CBO9780511840425> doi: 10.1017/cbo9780511840425
- Radford, A. (1990). *Syntactic theory and the acquisition of english syntax: The nature of early child grammars of english*. Oxford: Wiley Blackwell.
- Salhan, S., Diehl Martinez, R., Goriely, Z., & Buttery, P. (2024, November). Less is more: Pre-training cross-lingual small-scale language models with cognitively-plausible curriculum learning strategies. In M. Y. Hu et al. (Eds.), *The 2nd babylm challenge at the 28th conference on computational natural language learning* (pp. 174–188). Mi-

- ami, FL, USA: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/2024.conll-babylm.15/>
- Tomasello, M. (2000, March). Do young children have adult syntactic competence? *Cognition*, 74(3), 209–253. Retrieved from [http://dx.doi.org/10.1016/S0010-0277\(99\)00069-4](http://dx.doi.org/10.1016/S0010-0277(99)00069-4) doi: 10.1016/S0010-0277(99)00069-4
- Tomasello, M. (2005). *Constructing a language: A usage-based theory of language acquisition*. Harvard University Press. Retrieved from <http://dx.doi.org/10.2307/j.ctv26070v8> doi: 10.2307/j.ctv26070v8
- Westergaard, M. (2009). *The acquisition of word order: Micro-cues, information structure, and economy*. John Benjamins Publishing Company. Retrieved from <http://dx.doi.org/10.1075/la.145> doi: 10.1075/la.145