

An MG Parsing View into the Processing of Subject and Object Relative Clauses in Basque

Matteo Fiorini

University of Utah

matteo.fiorini@utah.edu

Jillian Chang

Great Neck South High School

jillianchang15@gmail.com

Aniello De Santo

University of Utah

aniello.desanto@utah.edu

Abstract

Stabler (2013)'s top-down parser for Minimalist grammars has been used to account for a variety of off-line processing preferences, with measures of memory load sensitive to subtle structural details. This paper expands the model's empirical coverage to ergative languages by looking at the processing asymmetries reported for Basque relative clauses. Our results show that the model predicts a subject over object preference as identified in the relevant psycholinguistic literature.

1 Introduction

A core question in research on human sentence processing is how language-specific linguistic features interact with more general processing mechanisms to give rise to the behavioral patterns recorded in production/comprehension experiments.

In this sense, differences between subject and object relative clauses (SRC and ORC, respectively) have received a lot of attention throughout the years (see Lau and Tanaka, 2021, for a review). Generally, SRCs are more common cross-linguistically (Keenan and Comrie, 1977), and they are reportedly produced and comprehended earlier and more easily than ORCs. While this subject advantage can be modulated by other properties of the sentence (e.g. case mismatches), it seems to be an overall strong pattern in both head-initial nominative/accusative languages with postnominal RCs (e.g., English or French; Mecklinger et al., 1995; Gibson, 1998; Frazier, 1987; Friedmann and Novogrodsky, 2004) and (somewhat less reliably) in head-final languages with prenominal RCs (e.g. Korean or Japanese; Kwon et al., 2010, 2013; Nakamura and Miyamoto, 2013).

Crucially, the very broad question about the interaction between language-specific properties and general cognitive processes of the human parser also leads to the more specialized question of *which* features of a language matter for different aspects

of sentence processing, and how. In particular, from the perspective of highly detailed syntactic frameworks, it seems important to probe the relevance of fine-grained syntactic details in deriving behavioral patterns (Miller and Chomsky, 1963; Bresnan, 1978; Rambow and Joshi, 1997).

In this paper, we follow work recasting this question in computational terms, by specifying a transparent linking hypothesis between the syntactic structures assumed in Minimalism (Chomsky, 1995) and off-line processing difficulty. Specifically, we adopt a model integrating Stabler (2013)'s top-down parser for Minimalist grammars (Stabler, 1996, 2011) with complexity metrics measuring memory usage to derive off-line estimates of processing complexity, based on the interaction between the parser's tree-traversal strategy and the rich structure of a derivation (Kobele et al., 2013; Gerth, 2015; Graf et al., 2017; De Santo, 2020b).

RCs in general, and the asymmetries in processing between subject and object RCs in particular, have been extensively probed with this model across a variety of languages (Graf et al., 2015, 2017; De Santo, 2021a,b; Zhang, 2017). Here then, we contribute to this line of work by evaluating the model's ability to predict the contrast between SRCs and ORCs reported for Basque. Basque is of particular interest to this type of investigation as a highly inflected, ergative, SOV language with both prenominal relatives and postnominal RCs. Ergative languages have been somewhat generally overlooked in past psycholinguistic work on the comprehension and production of RCs (Carreiras et al., 2010; Juncal Gutierrez Mangado and José Ezeizabarrena, 2012; Yetano and Laka, 2019, a.o.), as well as in the recent MG parsing literature. Their morpho-syntactic properties, however, make them ideal candidates to explore how the properties of RCs interact with processing strategies, as they pose challenges for various syntactic accounts of sentence structure *and* theories of sentence process-

ing proposed for other languages.

Due to its particular structural properties, Basque thus presents a novel, challenging test case for the computational model adopted in this paper. In showing how the model handles the SRC vs. ORC contrast reported by some of the Basque literature on RC comprehension, we not only extend the typological coverage of the model, but also highlight the relevance of computational models grounded in theoretical considerations in opening new research directions at the intersection between theoretical syntax and sentence processing.

2 Preliminaries: MG Parsing

Minimalist grammars (MGs; [Stabler, 1996, 2011](#)) are a lexicalized, feature driven formalism incorporating the structurally rich analysis of Minimalist syntax ([Chomsky, 1995, a.o.](#)).

An MG is a set of lexical items (LIs) consisting of a phonetic form and a finite, non-empty string of features. The latter are divided into two types: *Merge* features and *Move* features. LIs are assembled via the two relative feature checking operations *Merge* and *Move*. Essentially, Merge encodes subcategorization, while Move long-distance displacement dependencies. Given the scope of this paper, the technical details of the mechanism behind feature-checking are unnecessary — and in fact, in the rest of the paper we avoid displaying the feature component of the LIs altogether. What we want to highlight instead is the intuition behind the core MG data structure: *derivation trees*.

Intuitively, MG derivation trees encode the sequence of operations (Merge and Move) required to build the phrase structure tree for a specific sentence ([Michaelis, 1998](#); [Harkema, 2001](#); [Kobele et al., 2007](#)). Observe the tree in [Figure 1b](#), representing a simplified derivation of the sentence *Who does Salem like?*. Here, leaf nodes are labeled by LIs, while unary and binary branching nodes represent Move and Merge operations, respectively. The main and crucial difference between this representation and a more standard phrase structure tree is that in these derivations, moving phrases remain in their base position: their landing site can be fully (deterministically) reconstructed via feature calculus. What this means though is that the final word order of a sentence is not directly reflected in the order of the leaf nodes of a derivation tree. For the sake of clarity, while movement arrows are not technically part of this representation, since

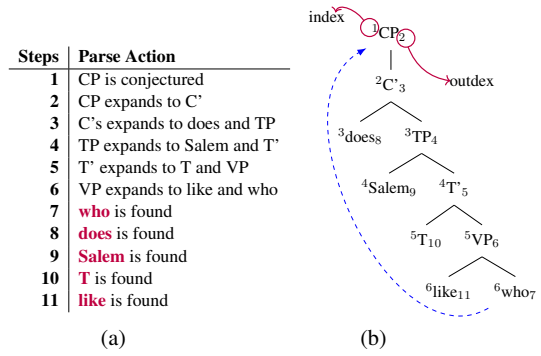


Figure 1: Example of a string-driven top-down tree traversal for an MG derivation tree of the sentence *Who does Salem like?*.

we make away with features in the rest of the paper, we will incorporate dashed arrows to indicate movement relations.

2.1 Top-Down Parsing

[Stabler \(2013\)](#) takes advantage of the fact that — modulo a more complex mapping from trees to strings — MG derivation trees form a regular tree language, to propose a string-driven MG variant of a standard depth-first, top-down parser for Context-Free Grammars. Essentially, this parser hypothesises tree nodes from top to bottom and from left to right. However, since the surface order of lexical items in the derivation tree is not the phrase structure tree’s surface order, simple left-to-right scanning of the leaf nodes yields the wrong order. The MG parser, while scanning the nodes, must thus also keep tracking the derivational operations which affect the linear word order and prioritizes resolving movement dependencies over the top-down strategy (i.e. the string-driven component).

Following [Kobele et al. \(2013\)](#), without delving too much in technical details, the parsing procedure can be outlined as follows: I) hypothesize the top of structure and add nodes downward (toward words) and left-to-right; II) if move is predicted, it triggers the search for mover → build the shortest path towards predicted mover; III) once the mover has been found, continue from the point where it was predicted. A memory stack plays a fundamental role in this: if a node is hypothesized at step i , but cannot be worked on until step j , it must be stored for $j - i$ steps in a priority queue.

The example in [Figure 1a](#) exemplifies this strategy for the tree in [Figure 1b](#). To keep track of these operations, we follow past literature on this topic

and adopt [Kobele et al. \(2013\)](#)’s notation: each node in the tree is annotated with a superscript (index) and a subscript (outdex). The annotation intuitively indicates for each node in the tree I) when it is first conjectured by the parser (index) and placed in the memory stack, and II) at what point it is considered completed and flushed from memory (outdex).

Since MGs are able to closely encode the detailed structural analyses of Minimalist syntax, [Stabler’s](#) MG parser has led to a rich line of work aimed at connecting syntactic assumptions to offline processing behavior, through the use of *complexity* metrics ([Kobele et al., 2013](#); [Gerth, 2015](#); [De Santo, 2020b](#), a.o.).

2.2 Complexity Metrics

We employ complexity metrics that predict processing difficulty based on how memory usage is affected by the geometry of the trees built by the parser.

Building on previous work in (computational) psycholinguistics ([Gibson, 1998](#); [Rambow and Joshi, 1997](#), a.o.), [Kobele et al. \(2013\)](#) identify broad cognitive notions of memory usage like 1) *tenure*: how long a node is kept in memory and 2) *size*: the amount of information a node consumes in memory. In practical terms, the *tenure* of a node is equal to the difference between its index and its outdex. Given how derivation trees are built by the parser, given a left-to-right string to tree matching a tenure of two is the minimum expected for the right sister in a tree with binary branching — thus, $\text{tenure} \leq 2$ is labelled as *trivial*. In practice, size encodes how nodes in a derivation consume memory because a phrase m moves across these nodes — and it can be computed in our simplified representation of derivation trees by subtracting the index of a moved element from the index of its landing site (see [Graf and Marcinek, 2014](#); [Graf et al., 2015](#), for a more technical discussion). For instance, referring to the annotated tree in [Figure 1b](#), the size of *who* is 6.

These memory notions can then be used to define a large set of complexity metrics measuring the offline processing difficulty over a full derivation tree. [Kobele et al. \(2007\)](#) associate tenure with quantitative values by implementing complexity metrics such as: $\text{MAXT} := \max(\text{tenure-of}(n))$, and $\text{SUMT} := \sum_n \text{tenure-of}(n)$. MAXT measures the maximum amount of time any node stays in mem-

ory during processing, while SUMT measures the overall amount of memory usage for all nodes with non-trivial tenure (i.e., > 2), capturing the total memory usage over the course of a parse. Building on these findings, [Graf and Marcinek \(2014\)](#) show that MAXT (only considering pronounced nodes) makes the right difficulty predictions for several phenomena, e.g., right vs. center embedding, nested vs. crossing dependencies, and the contrasts involving relative clauses at the center of our paper.

Following up on these results, [Graf et al. \(2015\)](#) extend the definition of these complexity measure to size. For instance, SUMSIZE can be used to measure the overall cost of maintaining long-distance filler-gap dependencies over a derivation. Let M be the set of all nodes of a derivation tree t that are the root of a subtree undergoing movement. For each $m \in M$, $i(m)$ is the index of m and $f(m)$ is the index of the highest Move node that m ’s subtree is moved to. Then SUMSIZE is defined as $\sum_{m \in M} i(m) - f(m)$.

[Graf et al. \(2015\)](#) also propose an idea similar to Optimality Theory’s ([Prince and Smolensky, 2008](#)) constraint ranking. In their formulation, metrics of the type $\langle M_1, M_2, \dots, M_n \rangle$ are ranked, and lower ranked metrics only matter if all higher ranked metrics have failed to pick out a unique winner (e.g., two constructions result in a *tie* over MAXT). While such a system would easily generate an enormous number of possible metrics, [Graf et al. \(2017\)](#) have argued that a small number of such metrics is in fact enough to account for a vast number of processing contrasts cross-linguistically. In particular, the ranking $\langle \text{MAXT}, \text{SUMSIZE} \rangle$ has been surprisingly successful in offering insights into the connection between processing load and syntactic choices (see [De Santo, 2020a](#); [Liu, 2018](#); [Lee, 2018](#); [De Santo and Lee, 2022](#), for additional support to these claims).

With a slight over-generalization of terminology, henceforth we refer to the combination of MGs as a grammar formalism, [Stabler \(2013\)](#)’s top-down parser, and complexity metrics estimating memory usage as the *MG Parser* or the *MG Model*, even though it is important to recognize that alternative combinations of these components are possible ([Yun et al., 2015](#); [Hunter, 2019](#); [Hunter et al., 2019](#)). Following [Graf et al. \(2017\)](#) and others, we then use this model to conduct pairwise comparisons of full derivations for constructions under analysis

(e.g., SRCs vs. ORCs) and derive estimates of processing difficulty that we can categorically match to the contrasts reported in the psycholinguistics literature.¹

3 SRCs vs ORCs in Basque

Ergative languages, albeit representing roughly 25% of the world languages (Dixon, 1994), have received relatively little attention in computational psycholinguistics' literature. As mentioned, Basque is an ergative and head-final language allowing for prenominal RCs (de Rijk, 2007). Furthermore, Basque is a three-way pro-drop language that can omit all arguments in a sentence (i.e., the XPs marked by ergative, absolutive, and dative case). Finally, while canonically an SOV language, the word order of Basque is prone to variation (de Rijk, 2007).

Importantly, while prenominal RCs have been modelled with the MG parser (Graf et al., 2017; Zhang, 2017) the focus (somewhat in parallel with the broader psycholinguistics literature) has been on East Asian languages (Japanese, Korean, and Mandarin Chinese more specifically). The availability of prenominal RCs combined with a highly flexible word order, and ergativity, however, makes Basque an ideal candidate to expand the array of languages the MG model has been tested against.

Consider now sentences like in (1), illustrating Basque's prenominal RC constructions. This example presents an SRC (1-a) and an ORC (1-b) in a subject-modifying set-up (that is, the RC modifies a noun acting as the subject of the main clause).

- (1) a. Irakasleak aipatu ditu-en
 teacher.PL.ABS mention.PRT has=comp
 ikasleak lagunak ditu
 student.SG.ERG friend.PL.ABS has
 orain.
 now
 'The student that mentioned the teachers has friends now.' **SRC**

- b. Irakasleak aipatu ditu-en
 teacher.PL.ABS mention.PRT has=comp
 ikasleak lagunak dira
 student.PL.ERG friend.PL.ABS are
 orain.
 now
 'The students that the teachers mentioned are friends now.' **ORC**

Consistently with other languages with prenominal RCs, behavioral experiments on Basque RCs preferences are somewhat split (cf. Kwon et al., 2013; Yang et al., 2010; Gibson and Wu, 2013; Yetano and Laka, 2019). However, there seems to be sound evidence that, in absence of other confounds (e.g. morphological and syntactic ambiguity) Basque participants show a clear subject preference (Juncal Gutierrez Mangado and José Ezeizabarrena, 2012; Munarriz et al., 2016; Yetano and Laka, 2019).

Additionally, recent studies on Basque have also shown a strong subject preference for postnominal RCs, a construction that seems to lack some of the morpho-syntactic ambiguity present in the prenominal structure (Carreiras et al., 2010; Yetano and Laka, 2019, a.o.). However, the syntactic status of postnominal RCs in Basque is controversial and understudied to the point that we are not aware of extensive theoretical work discussing the structural details of such configuration. Since syntactic choices are fundamental to the modelling approach taken here, in the present paper we leave postnominal structures aside, and focus on evaluating whether the parser can predict a subject advantage for prenominal sentences as in (1).

4 Modelling Basque RCs

As input to the MG parser we used derivations for the sentences in (1), as shown in Figure 2 and Figure 3. We expect a preference for SRCs over ORCs (SRC > ORC), following the results in (Juncal Gutierrez Mangado and José Ezeizabarrena, 2012; Munarriz et al., 2016). As the MG parser is sensitive to fine-grained structural details, we are interested in a) capturing current Minimalist approaches to the structure of these sentences and b) explore how much particular syntactic choices involved in the derivation of RCs affect the parser's prediction. Thus, we compare derivations following two different approaches to the structure of restrictive relatives. Note also that, given the combination of an SOV base-clause plus a prenominal RC construction, these sentences show a *gap-filler*

¹It is worth mentioning that in its full formulation, Stabler's parser exploits a search beam discarding the most unlikely predictions at each step. However, past work (starting with Kobele et al., 2013) has ignored the beam, assuming that the parser is equipped with a perfect oracle, which always makes the right choices when constructing a tree. On the one hand, such an idealization is obviously implausible from a psycholinguistics perspective. On the other hand, this choice allows the model to focus on the specific contribution of structure building operations to processing difficulty. Interestingly, even in this configuration the MG Parser has been used to gain insights into phenomena dealing with ambiguity resolution (De Santo and Lee, 2022; Lee and De Santo, 2022).




Figure 3: Annotated MG trees for the SRC and ORC sentences in Example (1) following the Head Internal analysis. Boxed nodes are those with tenure greater than 2.

tically (Bianchi, 2002a,b; Hemforth et al., 2015; Fernández, 2017, i.a.).

In particular, a number of assumptions are made in the proposed syntactic accounts in terms of functional projections and movement operations. Specifically, the HI analysis strongly relies on projections in the expanded *left periphery* (Rizzi, 1997) of the RC clause to derive the correct prenominal SOV surface order. These additional projections split the CP head into multiple projections, which encode aspectual and discourse-oriented information. Conversely, the HE analysis does not exploit these projections in the derivation, partially due to the fact that the modified noun is base-generated externally to the RC structure. It will thus be interesting to see whether and how these additional structural elements in the HI approach have any effect on the predictions made by the MG parser, when interacting with subject or object movement.

4.2 Modelling Results

As mentioned before, here we exclusively considered prenominal RCs like (1), and expect a SRC > ORC contrast. The (annotated) MG derivations for the two analyses are given in Figure 2 and Figure 3 respectively, for the prototypical sentences in (1).

With these preliminaries in place, we can look at the modelling results.² We evaluated the whole set of 1800 metrics defined by Graf et al. (2017) but, following previous MG parsing work, we focus our discussion on the predictions made by MAXT and SUMSIZE. As it turns out, the MG parser equipped with the $\langle \text{MAXT}, \text{SUMSIZE} \rangle$ metric predicts a preference for SRCs over ORCs for both analyses, but with interesting differences between the two in how this is accomplished (see Table 1). Note that for both derivations, the pairwise contrasts predicted

²All simulations were run with a version of the code made freely available by Graf et al. (2017) at <https://github.com/CompLab-StonyBrook/mgproc>.

	Head Internal			Head External		
	MaxT	Node	SumSize	MaxT	Node	SumSize
SRC	30	orain	50	25	orain	31
ORC	31	orain	64	25	orain	37

Table 1: Performance of MAXT and SUMSIZE for each of the RC sentences in (1), derived according to a Head Internal and a Head External analysis.

do not change whether we consider intermediate movement steps or not (cf. Zhang, 2017).

Consider again the sentences in (1). For the HE analysis, MAXT leads to a tie between SRC and ORC, with a tenure of 25 recorded on the matrix clause temporal auxiliary *orain*. Since we are considering subject RCs (i.e. the noun modified by the SRC/ORC goes to become the subject of the matrix clause) and because of the prenominal nature of the RC, every element in the matrix *vP* has to wait until the NP containing the RC and its noun moves to subject position in matrix Spec,TP. This results in an equivalent tenure on those nodes, given that the size (in terms of number of nodes) of the two structures is the same, independently of whether the head noun originates in subject or object position within the RC. Note that this tie is also shown on the rightmost node internal to the RC (*en*), illustrating how this is more a consequence of the prenominal RC than of having picked (consistently with the psycholinguistic literature) subject modifying structures. Interestingly, tenure on the head itself does display a subject preference. In the ORC case, *irakasleak* (in the embedded subject position) comes early in the linear sequence but has to wait for the the movement of OP from object position to Spec,CP to be resolved before it can be flushed out of the stack-memory of the parser. Nonetheless, the tie on MAXT is not a problem for a model using a ranked metric, and SUMSIZE makes in fact the correct prediction by capitalizing on the longer movement of OP in ORCs *and* on the additional movement of the embedded subject to (RC internal) Spec,TP.

Conversely, for the HI analysis, MAXT makes directly the correct prediction, registering a slightly higher tenure on the highest temporal adjunct (*orain*, but also on lower nodes) in the ORC structure. Inspecting the HI derivations more closely, we note that in this case tenure on the relativized head (*ikasleak*) predicts the opposite preference (24 – 16 for the SRC compared to a 25 – 20 for the ORC). This is due to the fact that in the SRC

construction *ikasleak* is predicted in Spec,*vP* but then it cannot be confirmed and discarded from memory until the lower VP elements (preceding it in the linear order) are found lower in the clause, and that their movement dependencies are resolved. Being predicted in object position makes it so that the waiting time for *ikasleak* is actually lower in the ORC derivation. Interestingly however, this difference disappears when we move to the higher parts of both derivations, covered by the movement of the head to Spec, ForP. The ORC derivation needing the additional movement of the RC-internal TP clause to Spec,DP is what causes its tenure to increase on the higher nodes compared to that for the SRC. SUMSIZE makes again the correct prediction by considering both these additional movement dependencies and the number of extra-projections the object head needs to move across compared to the subject head.

5 Discussion

The results above display how the MG model is able to predict a subject preference in Basque SRC/ORC constructions, in line with what reported in both production and comprehension studies (Juncal Gutierrez Mangado and José Ezeizabarrena, 2012; Munarriz et al., 2016). This success adds to previous MG modelling of sentence processing results in supporting MAXT and SUMSIZE as a combination of metrics able to capture different aspects of syntactic difficulty cross-linguistically, in ways that can give us insights into the relation between parsing and fine-grained syntactic choices.

Importantly, while the model predicts the SRC > ORC ranking across two different syntactic analyses of RCs, a closer inspection reveals that it does so in strikingly different ways. In this sense, the results for the HI analysis seem mostly driven by the additional structural operations required by that analysis to derive the correct linearization. In contrast, a study of the metrics’ values for the HE analysis show a higher sensitivity for the differences between subject and object RCs both in

terms of movement dependencies, and the way the tree traversal strategy of the MG parser interacts with subtler differences in the geometry of the two derivation trees. These considerations thus highlight the value of a model quantifying the relation between syntactic structure and processing load as transparently as possible, so to allow not just for quantitative predictions but also careful qualitative analyses. Specifically, this suggests ways in which this type of model could be used by both syntacticians and psycholinguists to spell out which aspects of a syntactic derivation they predict to be relevant to behavioural performance, and why.

Going back to the question of RC processing more broadly, past psycholinguistic literature has focused on well-established asymmetries between SRCs and ORCs in order to investigate the connection between universal properties of the human parser and the syntactic features of particular languages. In this sense, even though the MG model does not encode a bias towards structural locality explicitly, these results (together with previous MG modelling work on RC asymmetries in other languages) show how a subject preference could arise cross-linguistically from the interaction of language specific structural properties and generalist parsing mechanisms taking memory usage into consideration.

Finally, it is worth mentioning again that some experimental studies have reported a preference for an ORC interpretation in the processing of Basque prenominal RCs (Carreiras et al., 2010; Yetano and Laka, 2019). However, a close look at the kind of sentences tested in these studies has highlighted how the syntactic properties of Basque make prenominal SRCs temporarily ambiguous. Recall that Basque can drop several arguments (bearing ergative, absolutive, and indirect case). Additionally, the prenominal RC does not contain an explicit particle (like a *wh* element in English) functioning as a complementizer, which is instead attached to the subordinate verb in clause-final position (*en* in our sentences). Taken together, these characteristics make it so that a prenominal SRC could be initially interpreted as a main clause with dropped argument, at least until the parser reaches the embedded verb marked with the complementizer (Carreiras et al., 2010). Thus, past work has argued that the ORC preference found by some studies is in fact a result of the additional complexity brought in SRC structures by ambiguity resolution

(Juncal Gutierrez Mangado and José Ezeizabarrena, 2012). This explanation is also in line with what has been argued for the object preference sometimes found in other head-final languages (Kwon et al., 2010, 2013; Nakamura and Miyamoto, 2013). In fact, when testing unambiguous postnominal RCs (as in (2)), Yetano and Laka (2019) report a strong preference for SRC constructions.

- (2) a. Ikasle-a-ki, [zein-a-ki ei
Student-sg-erg, [who-sg-erg ei
irakasle-ak aipatu bait-ditu,]
teacher-pl mentioned Comp-has,]
lagun-ak ditu orain.
friend-pl has now.
'The student, who mentioned the teachers, has friends now.' **SRC**
- b. Ikasle-aki, [zein-aki irakasle-a-k ei
Student-pl, [who-pl teacher-sg-erg ei
aipatu bait-ditu,] lagun-ak dira
mentioned Comp-has,] friend-pl are
orain.
now.
'The students, who mentioned the teacher, are friends now.' **ORC**

As discussed before, here we did not test our model against postnominal structures, in part due to the lack of extensive syntactic literature on the topic. Importantly though, future work in this direction will have to consider the variety of morpho-syntactic factors (especially related to case syncreticity) differentiating prenominal and postnominal constructions, and suggest fundamental ways in which the current MG model (only sensitive to tree geometry) should be expanded in order to pursue a full, in-depth investigation of syntactic processing in ergative languages. On the other hand, the preliminary results in this paper draw attention to gaps in the literature connecting theoretical syntax and psycholinguistic studies, thus showcasing once again the contribution of models like the MG parser to the broader study of the role of syntactic representation in linguistic cognition.

Acknowledgments

The authors would like to thank the audience at HSP 2023, as well as the anonymous SCiL reviewers, for valuable feedback on this work.

References

- Xabier Artiagoitia. 1992. Why basque doesn't relative everything? *Anuario del Seminario de Filología Vasca "Julio de Urquijo"*, pages 11–35.
- Valentina Bianchi. 2002a. Headed relative clauses in generative syntax. Part I. *Glott International*, 6.7.
- Valentina Bianchi. 2002b. Headed relative clauses in generative syntax. Part II. *Glott International*, 6.8.
- Joan Bresnan. 1978. A realistic transformational grammar. In Morris Halle, Joan Bresnan, and George A. Miller, editors, *Linguistic Theory and Psychological Reality*, pages 1–59. The MIT Press.
- Manuel Carreiras, Jon Andoni Duñabeitia, Marta Vergara, Irene de la Cruz-Pavía, and Itziar Laka. 2010. Subject relative clauses are not universally easier to process: Evidence from basque. *Cognition*, 115(1):79–92.
- Noam A. Chomsky. 1995. *The Minimalist Program*. The MIT Press, Cambridge, MA.
- Rudolf P.G. de Rijk. 2007. *Standard Basque: A Progressive Grammar*. The MIT Press.
- Aniello De Santo. 2020a. MG parsing as a model of gradient acceptability in syntactic islands. In *Proceedings of the Society for Computation in Linguistics 2020*, pages 59–69.
- Aniello De Santo. 2020b. *Structure and memory: A computational model of storage, gradience, and priming*. Ph.D. thesis, State University of New York at Stony Brook.
- Aniello De Santo. 2021a. Italian postverbal subjects from a Minimalist parsing perspective. *Lingue e linguaggio*, 20(2):199–227.
- Aniello De Santo. 2021b. A Minimalist approach to facilitatory effects in stacked relative clauses. *Proceedings of the Society for Computation in Linguistics*, 4(1):1–17.
- Aniello De Santo and So Young Lee. 2022. [Evaluating structural economy claims in relative clause attachment](#). In *Proceedings of the Society for Computation in Linguistics 2022*, pages 65–75, online. Association for Computational Linguistics.
- Robert M.V. Dixon. 1994. *Ergativity*. Cambridge University Press.
- Eva M. Fernández. 2017. The prosody produced by spanish-english bilinguals: A preliminary investigations and implications for sentence processing.
- Lyn Frazier. 1987. Syntactic processing: evidence from dutch. *Natural Language & Linguistic Theory*, pages 519–559.
- Naama Friedmann and Rama Novogrodsky. 2004. The acquisition of relative clause comprehension in hebrew: A study of sli and normal development. *Journal of Child language*, 31(3):661–681.
- Sabrina Gerth. 2015. Memory limitations in sentence comprehension: a structural-based complexity metric of processing difficulty. Universitätsverlag Potsdam.
- Edward Gibson. 1998. Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1):1–76.
- Edward Gibson and H-H Iris Wu. 2013. Processing chinese relative clauses in context. *Language and Cognitive Processes*, 28(1-2):125–155.
- Ager Gondra. 2015. [Head raising analysis and case reevaluation](#). *Borealis : An International Journal of Hispanic Linguistics*, 4:193.
- Ager Gondra. 2016a. *Basque relative clauses: Head raising, case and micro-variation within Bizkaiera*. Purdue University, PhD Dissertation.
- Ager Gondra. 2016b. Universal grammar: Multiple strategies to construct relative clauses. *Estudios interlingüísticos*, (4):59–75.
- Thomas Graf, Brigitta Fodor, James Monette, Gianpaul Rachiele, Aunika Warren, and Chong Zhang. 2015. [A refined notion of memory usage for minimalist parsing](#). In *Proceedings of the 14th Meeting on the Mathematics of Language (MoL 2015)*, pages 1–14, Chicago, USA. Association for Computational Linguistics.
- Thomas Graf and Bradley Marcinek. 2014. [Evaluating evaluation metrics for minimalist parsing](#). In *Proceedings of the Fifth Workshop on Cognitive Modeling and Computational Linguistics*, pages 28–36, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Thomas Graf, James Monette, and Chong Zhang. 2017. [Relative clauses as a benchmark for minimalist parsing](#). *Journal of Language Modelling*, 5(1):57–106.
- Henk Harkema. 2001. A characterization of minimalist languages. In *International Conference on Logical Aspects of Computational Linguistics*, pages 193–211. Springer.
- Barbara Hemforth, Susana Fernandez, Charles Clifton, Lyn Frazier, Lars Konieczny, and Michael Walter. 2015. [Relative clause attachment in german, english, spanish and french: Effects of position and length](#). *Lingua*, 166:43–64.
- Tim Hunter. 2019. Left-corner parsing of Minimalist Grammars. *Minimalist Parsing*.
- Tim Hunter, Miloš Stanojević, and Edward Stabler. 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*, pages 1–10.

- Maria Juncal Gutierrez Mangado and María José Ezeizabarrena. 2012. Asymmetry in child comprehension and production of basque subject and object relative clauses. In *BUCLD proceedings online*.
- Richard Kayne. 1994. *The antisymmetry of syntax*. MIT Press, Cambridge, MA.
- Edward L. Keenan and Bernard Comrie. 1977. Noun phrase accessibility and universal grammar. *Linguistic Inquiry*, 8(1):63–99.
- Gregory M. Kobele, Sabrina Gerth, and John Hale. 2013. Memory resource allocation in top-down minimalist parsing. In *Formal Grammar*, pages 32–51, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Gregory M. Kobele, Christian Retoré, and Sylvain Salvati. 2007. An automata-theoretic approach to minimalism. In *Model Theoretic Syntax at 10*, pages 73–82.
- Nayoung Kwon, Robert Kluender, Marta Kutas, and Maria Polinsky. 2013. Subject/object processing asymmetries in korean relative clauses: Evidence from erp data. *Language*, 89(3):537.
- Nayoung Kwon, Yoonhyoung Lee, Peter C Gordon, Robert Kluender, and Maria Polinsky. 2010. Cognitive and linguistic factors affecting subject/object asymmetry: An eye-tracking study of prenominal relative clauses in korean. *Language*, pages 546–582.
- Elaine Lau and Nozomi Tanaka. 2021. The subject advantage in relative clauses: A review. *Glossa: a journal of general linguistics*, 6(1).
- So Young Lee. 2018. [A minimalist parsing account of attachment ambiguity in english and korean](#). *Journal of Cognitive Science*, 19(3):291–329.
- So Young Lee and Aniello De Santo. 2022. A computational view into the structure of attachment ambiguities in Chinese and Korean. In *Proceedings of NELS 52*, online.
- Lei Liu. 2018. [Minimalist parsing of heavy NP shift](#). In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation*, Hong Kong. Association for Computational Linguistics.
- Axel Mecklinger, Herbert Schriefers, Karsten Steinhauer, and Angela D Friederici. 1995. Processing relative clauses varying on syntactic and semantic dimensions: An analysis with event-related potentials. *Memory & Cognition*, 23(4):477–494.
- Jens Michaelis. 1998. Derivational minimalism is mildly context-sensitive. In *International Conference on Logical Aspects of Computational Linguistics*, pages 179–198. Springer.
- George A. Miller and Noam Chomsky. 1963. Finitary models of language users. In D. Luce, editor, *Handbook of Mathematical Psychology*, pages 2–419. John Wiley & Sons.
- Amaia Munarriz, Maria-Jose Ezeizabarrena, and M JUNCAL GUTIERREZ-MANGADO. 2016. Differential and selective morpho-syntactic impairment in spanish-basque bilingual aphasia. *Bilingualism: Language and Cognition*, 19(4):810–833.
- Michiko Nakamura and Edson T Miyamoto. 2013. The object before subject bias and the processing of double-gap relative clauses in japanese. *Language and Cognitive Processes*, 28(3):303–334.
- Alan Prince and Paul Smolensky. 2008. *Optimality Theory: Constraint Interaction in Generative Grammar*, chapter 1. John Wiley Sons, Ltd.
- Owen Rambow and Aravind Joshi. 1997. A formal look at dependency grammars and phrase-structure grammars, with special consideration of word-order phenomena. *Recent trends in meaning-text theory*, 39:167–190.
- Luigi Rizzi. 1997. The fine structure of the left periphery. *Elements of Grammar: Handbook in Generative Syntax*, pages 281–337.
- Edward P. Stabler. 1996. Derivational minimalism. In *International Conference on Logical Aspects of Computational Linguistics*, pages 68–95. Springer.
- Edward P. Stabler. 2011. [617 Computational Perspectives on Minimalism](#). In *The Oxford Handbook of Linguistic Minimalism*. Oxford University Press.
- Edward P. Stabler. 2013. Two models of minimalist, incremental syntactic analysis. *Topics in cognitive science*, 5(3):611–633.
- Chin Lung Yang, Charles A Perfetti, and Ying Liu. 2010. Sentence integration processes: An erp study of chinese sentence comprehension with relative clauses. *Brain and Language*, 112(2):85–100.
- Iraia Yetano and Itziar Laka. 2019. Processing preferences in an ergative language: Evidence from basque postnominal relative clauses. *Pello Salabururi esker onez*, page 137.
- Jiwon Yun, Zhong Chen, Tim Hunter, John Whitman, and John Hale. 2015. Uncertainty in processing relative clauses across East Asian languages. *Journal of East Asian Linguistics*, 24:113–148.
- Chong Zhang. 2017. *Stacked relatives: their structure, processing and computation*. Ph.D. thesis, State University of New York at Stony Brook.