

Recursive numeral systems optimize the trade-off between lexicon size and average morphosyntactic complexity

Milica Denić¹ and Jakub Szymanik²

¹*Tel Aviv University*

²*University of Trento*

Abstract

Human languages vary in terms of which meanings they lexicalize, but there are important constraints on this variation. It has been argued that languages are under two competing pressures: the pressure to be simple (e.g., to have a small lexicon size) and to allow for an informative (i.e., precise) communication with their lexical items, and that which meanings get lexicalized may be explained by languages finding a good way to trade off between these two pressures (Kemp and Regier 2012 and much subsequent work). However, in certain semantic domains, it is possible to reach very high levels of informativeness even if very few meanings from that domain are lexicalized. This is due to productive morphosyntax, which may allow for construction of meanings which are not lexicalized. Consider the semantic domain of natural numbers: many languages lexicalize few natural number meanings as monomorphemic expressions, but can precisely convey any natural number meaning using morphosyntactically complex numerals. In such semantic domains, lexicon size is not in direct competition with informativeness. What explains which meanings are lexicalized in such semantic domains? We will argue that in such cases, languages are (near-)optimal solutions to a different kind of trade-off problem: the trade-off between the pressure to lexicalize as few meanings as possible (i.e, to minimize lexicon size) and the pressure to produce as morphosyntactically simple utterances as possible (i.e, to minimize average morphosyntactic complexity of utterances). This study in conjunction with previous work on communicative efficiency suggests that, in order to explain which meanings get lexicalized across languages and across semantic domains, a more general approach may be that languages are finding a good way to trade off between not two but *three* pressures: be simple, be informative, and minimize average morphosyntactic complexity of utterances.

Keywords: numerals; number; simplicity; informativeness; average morphosyntactic complexity; trade-off

1 Introduction

Human languages vary in terms of which meanings they lexicalize into simple morphemes.¹ There are nonetheless important constraints on this variation: some meanings or meaning types are very frequently lexicalized, while others very rarely. Such constraints have been identified in both content and functional vocabulary. For instance, in the domain of content words, color terms have a well-established order of appearance in language evolution — certain semantic contrasts must be established before others (Berlin and Kay 1969). In addition, color terms label only convex regions of conceptual color space (Gärdenfors 2014, Jäger 2010). In the domain of function words, multiple semantic constraints have been identified for quantificational determiners (Horn 1972, Barwise and Cooper 1981, Keenan and Stavi 1986, Peters and Westerståhl 2006, Hackl 2009). For instance, no language lexicalizes a quantificational determiner whose meaning is *not every* (Horn 1972). Where do these constraints come from? Three prominent answers to this questions include learnability (e.g., certain meanings are rarely or not lexicalized because they are harder to learn, cf. Hunter and Lidz 2013, Chemla, Buccola, and Dautriche 2019, Steinert-Threlkeld and Szymanik 2018, 2020, Maldonado and Culbertson 2022), syntax-semantics interface (e.g., certain meanings are not lexicalized because they would be degenerate due to how syntactic structures are semantically interpreted, cf. Fox 2002, Romoli 2015), and communicative efficiency (e.g., certain meanings are rarely or not lexicalized because lexicalizing them wouldn't improve communicative efficiency of a language, cf. Kemp and Regier 2012 and much subsequent work).

In this paper, we won't discuss learnability and syntax-semantics interface explanations, and will focus instead on the communicative efficiency explanation. According to the most prominent version of the communicative efficiency explanation, languages are under the pressure to be simple (e.g., to have a small lexicon) while simultaneously being under the pressure to be informative, i.e., to allow for a precise communication. These two pressures are in competition. For instance, if a language only has one word for colors, it will have a small lexicon, but it won't allow for a very precise communication about colors. Adding more words for various colors would allow for a more precise communication, but at the cost of having a larger lexicon. Natural languages have been argued to lexicalize the meanings which allow them to achieve a good compromise between these two pressures, i.e., *to optimize the simplicity/informativeness trade-off* (Kemp and Regier 2012, Regier, Kemp, and Kay 2015, Xu, Regier, and Malt 2016, Xu, Liu, and Regier 2020, Kemp, Xu, and Regier 2018, Zaslavsky, Kemp, Regier, and Tishby 2018, Steinert-Threlkeld 2019, 2021, Mollica, Bacon, Zaslavsky, Xu, Regier, and Kemp 2021, Denić, Steinert-Threlkeld, and Szymanik 2021, 2022, Zaslavsky, Maldonado, and Culbertson 2021, Uegaki 2022).

However, in certain semantic domains, it is possible to reach very high levels of informativeness even if very few meanings from that semantic domain are lexicalized. This is due to productive morphosyntax, which may allow for construction of meanings which are not lexicalized. Consider a semantic domain of natural numbers in languages with a so-called *recursive numeral system*: these are languages in which any natural number meaning can be expressed precisely (e.g., English). Many such languages lexicalize few

¹We consider that a meaning is lexicalized if there is a simple morpheme carrying this meaning in a language. Whether this morpheme can or can't stand alone as a single word is not part of our criterion for lexicalization. For instance, English lexicalizes the number meaning *six* into a morpheme which can stand alone as a single word, and it lexicalizes plural meaning in a morpheme *-s*, among others, which cannot stand alone as a single word.

natural number meanings, but can precisely convey any natural number meaning using morphosyntactically complex numerals (e.g., *sixty-one* in English). In such semantic domains, lexicon size is not in direct competition with informativeness. What explains which meanings are lexicalized in these domains?

We will propose that, in semantic domains in which productive morphosyntax enables precise communication even with very few lexicalized meanings, languages are under the pressure to lexicalize as few meanings as possible (i.e, to minimize lexicon size) and the pressure to produce as morphosyntactically simple utterances as possible (i.e, to minimize average morphosyntactic complexity of utterances). Lexicon size and average morphosyntactic complexity of utterances are in competition in such domains: reducing average morphosyntactic complexity of utterances will often require lexicalizing more meanings, and reducing the size of the lexicon will often result in needing utterances of greater morphosyntactic complexity to communicate. *We thus propose that, in such domains, languages lexicalize those meanings which allow them to optimize the trade-off between lexicon size and average morphosyntactic complexity of utterances.*

We will evaluate this proposal within a case study on lexicalized number meanings in languages with a recursive numeral system, and present evidence that recursive numeral systems indeed optimize the trade-off between lexicon size and average morphosyntactic complexity of numerals. This conclusion is in tension with [Xu et al. \(2020\)](#), who have instead argued that numeral systems—including recursive numeral systems—optimize the simplicity/informativeness trade-off. We review their results in Section 5; we will argue that their results show that other types of numeral systems optimize the simplicity/informativeness trade-off, but that they do not show that recursive numeral systems optimize the simplicity/informativeness trade-off.

Once our results are in place, we will abstract away from numeral systems and discuss which notion of communicative efficiency may explain what meanings get lexicalized across semantic categories more generally, in light of this result and previous work within the simplicity/informativeness trade-off optimization approach (e.g., [Kemp and Regier 2012](#)). We will propose that, in light of the present case study on numeral systems *and* the previous work on simplicity/informativeness optimization approach, in order to explain which meanings get lexicalized across languages and across semantic domains, a more general approach may be that languages are finding a good compromise between not two but *three* pressures: be simple, be informative, and minimize average morphosyntactic complexity.

Finally, we will turn to several other lines of work ([Zipf 1949](#), [Piantadosi, Tily, and Gibson 2011](#), [Haspelmath 2021](#), [Mollica et al. 2021](#), [Carcassi and Sbardolini 2022](#)) which share with the present work the idea that speakers attempt to minimize complexity of their utterances, and explain in what way they are different from the proposal we are pursuing.

2 Experiment

We will present an experiment investigating how close natural languages’ recursive numeral systems are to trading off optimally between lexicon size and average morphosyntactic complexity of numerals. We will evaluate this by comparing natural languages to artificially generated languages. The artificial languages represent the space of possibilities in terms of lexicon size/average morphosyntactic complexity of numerals trade-off in

recursive numeral systems, and will reveal what the (approximately) optimal solutions to the trade-off problem are.

2.1 Natural languages

We assume that numerals across languages semantically denote numbers (e.g., the numeral *two* denotes the number 2), noting that this is a simplification (see [Bylinina and Nouwen 2020](#), [Spector 2013](#)). We collected cross-linguistic data on numerals denoting numbers 1-99 and their morphosyntactic components in recursive numeral systems from the sample of languages in the *Numeral bases* chapter in *The World Atlas of Language Structures (WALS)* [Comrie \(2013\)](#).² Out of 172 recursive numeral systems in [Comrie \(2013\)](#), 44 were excluded due to challenges with data collection or data interpretation.³ 128 languages were thus included in the analysis. *WALS* language samples are compiled with an aim to maximize genealogical and areal diversity of languages in them ([Comrie, Dryer, Gil, and Haspelmath 2013](#)) — we can thus have some confidence that we are analyzing a representative sample of world’s languages’ recursive numeral systems.

For each of the 128 studied numeral systems, for each numeral denoting a number in the range 1-99, its morphosyntactic components and their denotations were identified (cf. [Table 1](#) and [Table 2](#) for a few examples of numerals in Georgian and Fulfulde respectively). These numeral systems differ in terms of which number meanings they lexicalize. Some recurring options are listed in (1). The generalization seems to be that (in most cases) they lexicalize the following numbers from the range 1-99: (i) the first n numbers, with n varying across languages, often the first five or the first 10, and (ii) a couple of additional numbers such as 10 and/or 20 as in (1b) or (1c).

- (1) Lexicalized number concepts:
 - a. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 (74 languages)
 - b. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 20 (20 languages)
 - c. 1, 2, 3, 4, 5, 10 (9 languages)

Even when they lexicalize the same number meanings, languages sometimes make different choices in constructing morphosyntactically complex numerals. For instance, Kunama and Fulfulde lexicalize number meanings in (1c); Kunama constructs the numeral for number 9 as $10 - 1$, while Fulfulde does it as $5 + 4$. As another example, Greek and Georgian lexicalize number meanings in (1b); Greek constructs the numeral for number 45 as $4 \cdot 10 + 5$, while Georgian does it as $2 \cdot 20 + 5$. Morphosyntactically

²Two main sources were used to collect the cross-linguistic data. The primary source were descriptive grammars of individual languages, in most cases those referenced in [Comrie \(2013\)](#). When no descriptive grammar of a language was accessible to us, we used as a secondary source the data from the website <https://lingweb.eva.mpg.de/channumerals/>, maintained by Eugene Chen. This website is a collective effort of language scholars to document world’s language’s numeral systems. The list of analyzed languages, detailed descriptions of their numeral systems and the sources used for each language can be found in Appendix at: https://drive.google.com/file/d/1xZ8MfRDaPha_ggkU0iDJASGVe6HWSzi1/view?usp=sharing.

³For some of the languages from the sample in [Comrie \(2013\)](#), no complete description of the numeral system was accessible to us. Furthermore, a small number of languages were excluded due to difficulties with data interpretation. In most cases, this happened when morphosyntax of certain numerals was not aligned with their interpretation (e.g. in Zoque, the numeral for number 9 is morphologically 6+4; this is dubbed ‘correct misinterpretation’ in [Hurford \(1975\)](#)). Additionally, Danish was excluded for reasons explained in footnote 4 and Ainu was excluded for reasons explained in footnote 6.

Table 1: Georgian numerals for numbers 6, 30 and 62

Denotation (numeral)	Morphosyntactic make-up	Number of morphemes
6 (<i>ekvsi</i>)	6 (<i>ekvsi</i>)	1
30 (<i>otsdaati</i>)	20 (<i>ots-</i>) + (<i>-da-</i>) 10 (<i>-ati</i>)	3
62 (<i>samotsdaori</i>)	3 (<i>sam-</i>) · (\emptyset) 20 (<i>-ots-</i>) + (<i>-da-</i>) 2 (<i>-ori</i>)	5

complex numerals for numbers 1-99 across languages reveal that addition, multiplication and subtraction are productively involved in their construction⁴ (cf. Hurford 1975, 2007). Both number-denoting morphemes and morphemes denoting arithmetic operators can in principle be phonetically overt or covert: in practice, number-denoting morphemes are very rarely covert, while morphemes denoting arithmetic operators often are. Furthermore, certain number-denoting morphemes play a special role in the construction of morphosyntactically complex numerals (the so-called *bases*). For instance, in ‘base-10 languages’, morphosyntactically complex numerals for numbers up to 99 are in general constructed according to the morphosyntactic pattern $x \cdot 10 + n$ (e.g., Greek). On the other hand, in ‘base-20 languages’ morphosyntactically complex numerals for numbers up to 99 are in general constructed according to the morphosyntactic pattern $x \cdot 20 + n$ (e.g., Georgian). Many languages behave as ‘base-5’ languages when it comes to the composition of numerals for numbers 6-9, which they construct as $5 + n$ (e.g., Fulfulde).

Finding morphosyntactic components of numerals required studying the morphosyntactic patterns for numerals in each of the 128 languages. In most languages, the relevant morphosyntactic patterns were described in the descriptive grammars we consulted, or were otherwise easy to detect. There were however more difficult cases, even in languages we (the authors) are closely familiar with (for instance, is there a phonetic variant of the morpheme for number 2 in English *twelve*?). In such cases, we applied the following decision rule: if a numeral is an exception to an established morphosyntactic pattern in a language, but shows (phonetic or orthographic, depending on the available data) elements of morpheme(s) which should have been there if the morphosyntactic pattern was respected, we assume that the numeral follows the morphosyntactic pattern but with phonetic variants not seen elsewhere; otherwise we assume that the numeral is monomorphemic. For instance, English is a base-10 language and because of this we may expect that the numeral for 12 will be built from morphemes for 2 and 10; as the numeral *twelve* has elements of other phonetic realizations of the morpheme for the number 2 (e.g., *two*, *twe(n?)*- from *twenty*), we assume that English *twelve* incorporates the morpheme for number 2 (*twe-?*) and 10 (*-lve?*). Of course, one could argue for the application of an alternative decision rule (unless the morphosyntactic pattern is transparent, assume that the numeral is monomorphemic). It turns out that in practice applying this alternative decision wouldn’t qualitatively alter the results; cf. footnote 7.

⁴In the Danish numeral system, a morpheme denoting the fraction *half* is used to construct certain complex numerals; for instance the numeral fifty in Danish contains morphemes three, half and twenty ((three - half) · twenty). This suggests that there is a limited use of division too involved in the composition of numerals. As, to our knowledge, the extent to which division can be used productively in numeral systems is not well understood (cf. Hurford 1975, 2007), we assume for simplicity that division is not available, and exclude Danish from the corpus of natural languages.

Table 2: Fulfulde numerals for numbers 6, 30 and 62

Denotation (numeral)	Morphosyntactic make-up	Number of morphemes
6 (<i>jowe go'o</i>)	5 (<i>jowe</i>) + (\emptyset) 1 (<i>go'o</i>)	3
30 (<i>chappande tatti</i>)	10 (<i>chappande</i>) · (\emptyset) 3 (<i>tatti</i>)	3
62 (<i>chappande jowe go'o i didi</i>)	10 (<i>chappande</i>) · (\emptyset) (5 (<i>jowe</i>) + (\emptyset) 1 (<i>go'o</i>)) + (<i>i</i>) 2 (<i>didi</i>)	7

2.2 Lexicon size and average morphosyntactic complexity

We consider that lexicon size of a language is the number of lexicalized meanings — rather than the number of form-meaning pairs, e.g., we consider that English has one lexicon entry for 10 which can be phonetically realized in multiple ways (e.g., *ten*, *-teen*, *-ty*). This is important to keep in mind as it affects how our results should be interpreted: we will argue that languages optimize the trade-off between the number of lexicalized meanings and average morphosyntactic complexity of utterances, rather than the trade-off between the number of form-meaning pairs and average morphosyntactic complexity of utterances. Of course, an important question that remains open is why languages sometimes have multiple phonetic realizations of a single meaning.

Average morphosyntactic complexity of numerals in a language L is computed according to the formula in (2). In (2), $ms_complexity(n, L)$ is the morphosyntactic complexity of the numeral (i.e., the number of morphemes in it) of the language L denoting the number n and $P(n)$ is the probability that the number n needs to be communicated. For instance, average morphosyntactic complexity of English is obtained as $P(1) \times$ the number of morphemes in ‘one’ + $P(2) \times$ the number of morphemes in ‘two’ + ... + $P(99) \times$ the number of morphemes in ‘ninety-nine’. We assume that the probabilities that different numbers need to be communicated follow a power-law distribution as in (3) (cf. Dehaene and Mehler 1992, Piantadosi 2016, Xu et al. 2020).⁵ Qualitatively, this probability distribution captures that the larger the number n , the lower the need to talk about it.

(2)

$$average_ms_complexity(L) = \sum_{n \in [1, 99]} P(n) \cdot ms_complexity(n, L)$$

(3) **Prior over numbers:**

$$P(n) \propto n^{-2}$$

⁵The distribution in (3) is a simplification of a more complex cross-cultural reality. While previous work has established that numerals across languages generally follow such a distribution, there are also various deviations across languages (Dehaene and Mehler 1992). For instance, numerals for round numbers (e.g., numerals for 10, 20 in English) are used more frequently than expected given (3); to our knowledge, there is no consensus yet on what makes a number round and whether the same numbers are round across cultures. Furthermore, certain numbers may be referred to more or less frequently than expected given (3) in a language because of culture-specific reasons (e.g., unlucky number 13).

2.3 Approximating the Pareto frontier

In order to evaluate whether natural languages with a recursive numeral system optimize the trade-off between lexicon size and average morphosyntactic complexity of numerals, we need to establish how close they are to the optimal solutions to the lexicon size/average morphosyntactic complexity of numerals trade-off problem. The set of optimal solutions — called *the Pareto frontier* — is a set of (theoretically possible) languages for which there is no other (theoretically possible) language which is better on one of the two dimensions (lexicon size, average morphosyntactic complexity of numerals) without being worse on the other.

We use an evolutionary algorithm to approximate the Pareto frontier (cf. Steinert-Threlkeld 2019, 2021, Denić et al. 2021, 2022), which involves generations of many artificial languages (i.e., artificial recursive numeral systems). Before we explain how the evolutionary algorithm works, we explain how an artificial language is generated.

We use the grammar proposed by Hurford (1975, 2007) for natural language numerals generation to generate numerals of artificial languages. Hurford (1975, 2007) in his seminal work proposes that number-denoting morphemes across languages are of the syntactic category Digits (D) or Multipliers (M). Roughly, morphemes denoting numbers which are ‘bases’ as 10 or 20 in ‘base-10’ or ‘base-20’ languages are of the syntactic category M ; other morphemes are of the syntactic category D . He proposes that, across languages, numerals are constructed according to the grammar in (4). For instance, consider a language whose $D = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$, $M = \{10\}$ — in this language, $5 \cdot 10 + 6$, $5 \cdot 10 - 6$, $10 + 4$, $10 - 4$ are examples of morphosyntactically well-formed expressions, while $5 \cdot 2$, $5 + 3$ are examples of expressions which are not well-formed.⁶

- (4) NUMBER \longrightarrow D | PHRASE | PHRASE + NUMBER | PHRASE – NUMBER
PHRASE \longrightarrow M | NUMBER \cdot M

For our purposes, an artificial language consists of (a) a lexicon of number-denoting morphemes, which are of category D or M (cf. the grammar in (4)), and (b) a set of numerals for numbers 1-99 generated according to the morphosyntactic rules in (4) from the lexicon, where each numeral is the shortest expression (or one of them, in case of a tie) for a given number. For example, an artificial language L may have in its lexicon morphemes for 1, 2, 4, 6, 8, 10, such that 1, 2 and 4 are of category D and 6, 8 and 10 are of category M . Numerals of L for other numbers in range 1-99 would be generated from these morphemes as explained above. For instance, the numeral for number 16 in L would be randomly selected from one of the three morphosyntactically simplest options: $10 + 6$, $2 \cdot 8$, $8 + 8$. These can be generated by the grammar in (4) and are of equal morphosyntactic complexity (each has 3 morphemes). On the other hand, $4 \cdot 4$ is not an option for numeral 16 in L because, even though it would also have three morphemes, it cannot be generated by the grammar in (4) (because 4 is of category D in L). Similarly, $2 \cdot 6 + 4$ is not an option for number 16 in L because, even though it can be generated by the grammar in (4), it is not one of the morphosyntactically simplest options (it has

⁶ All numerals in all natural languages in this study can be constructed using the grammar in (4). One language in *WALS* corpus, Ainu, has numerals which cannot be derived by this grammar. Hurford (1975) is aware of this and discusses the case of Ainu at length; he considers different extensions of the grammar in (4) to accommodate Ainu but provides arguments against each of them, thus leaving the problem of Ainu unresolved. Because we will rely on the grammar in (4) to generate artificial languages and to estimate the Pareto frontier (cf. Section 2.3), we excluded Ainu from the corpus of natural languages.

5 morphemes).

Due to computational constraints, two restrictions are imposed on the search for the shortest expression for a given number in a language. (1) Expressions of depth x (i.e., expressions with at most x number-denoting morphemes and $x - 1$ arithmetic operators) are incrementally constructed from expressions of lower depths (e.g., expressions of depth 2 are constructed by combining expressions of depth 1). However, at all depths, the meaning of expressions is restricted to be a natural number in $[1, 200]$. (2) If no expression for a number meaning is found with a depth of at most 7, the search is abandoned, and the language is discarded.

The evolutionary algorithm works as follows. First, the generation 0 is created, which consists of 2000 artificial languages. The lexicons of these artificial languages are generated by drawing two random samples of numbers between 1-99; these stand for morphemes of category D and M respectively. As natural languages tend to have very few morphemes of category M (often only 1 or 2), we restrict the size of the random sample for the category M to at most 5 (no such restriction is imposed on the category D). The numerals for numbers 1-99 of these languages are generated as explained above. The dominant languages of a generation (those for which there is no language which is better on one dimension of lexicon size and average morphosyntactic complexity without being worse on the other) each give rise to an equal number of offspring languages, which are obtained via a small number of mutations (between 1 and 3; these mutations include removing a number morpheme from D or M , adding a number morpheme to D or M , and interchanging a number morpheme in D or M , making sure that M is no larger than 5) from dominant languages. The dominant languages from generation 0 together with their offspring languages constitute generation 1, whose size is limited to 2000 languages. This process is repeated for 100 generations. Finally, the dominant languages are selected from the union of the last generation and the natural languages. Each of these dominant languages is a point in a two-dimensional (lexicon size and average morphosyntactic complexity of numerals) space; we do spline interpolation of these points to form a Pareto frontier.

2.4 Results

The natural languages and the artificial languages generated through the 100 generations of the evolutionary algorithm are plotted in Figure 1. The approximated Pareto frontier — i.e., the set of (nearly-)optimal solutions to the problem of trading off between lexicon size and average morphosyntactic complexity of numerals — is plotted as the black curve in Figure 1.

We first note that natural languages all lie along or very close to the Pareto frontier in Figure 1. This speaks in favor of natural languages' recursive numeral systems optimizing the trade-off between lexicon size and average morphosyntactic complexity of numerals.⁷

⁷ Recall that, when the number of morphemes in a numeral was difficult to establish, we applied the following decision rule: if a numeral is an exception to an established morphosyntactic pattern in a language, but shows (phonetic or orthographic, depending on the available data) elements of morpheme(s) which should have been there if the morphosyntactic pattern was respected, we assume that the numeral follows the morphosyntactic pattern but with phonetic variants not seen elsewhere; otherwise we assume that the numeral is monomorphemic. The alternative would have been to treat all difficult cases as monomorphemic. Importantly, however, applying the alternative decision rule (treat all difficult cases as monomorphemic) wouldn't alter the results qualitatively. With such an alternative decision rule, natural languages would have slightly larger lexicon sizes and slightly lower average morphosyntactic

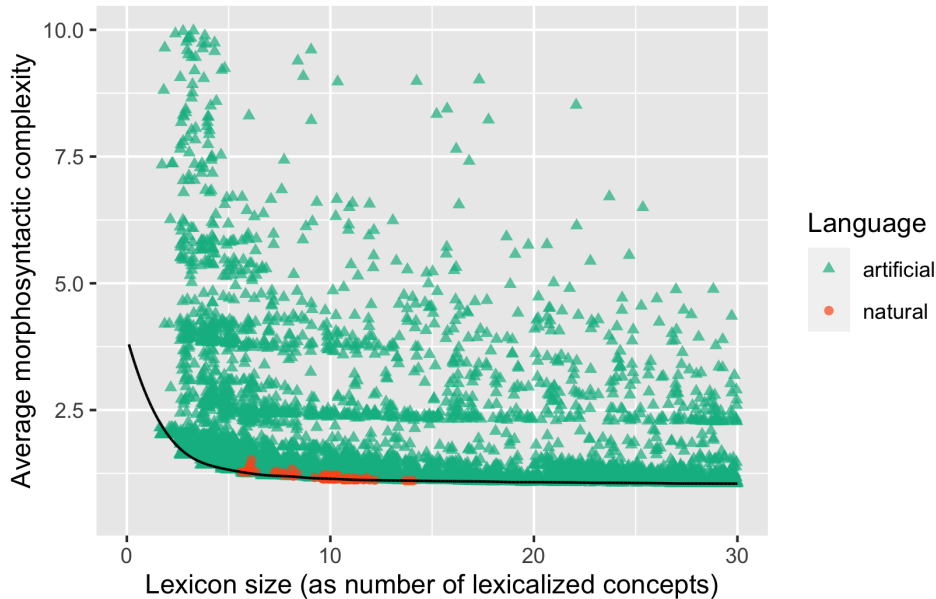


Figure 1: Experiment: Lexicon size and average morphosyntactic complexity of natural languages compared to artificial languages generated via an evolutionary algorithm. Natural languages lie at or very close to the Pareto frontier (black curve).

Interestingly, there is some variation in terms of where along the Pareto frontier natural languages lie. In other words, natural languages differ in terms of which optimal solution to the lexicon size/average morphosyntactic complexity of utterances trade-off problem they are approaching. Take for example cases of Georgian and Fulfulde: a few examples of their numerals are in Tables 1 and 2 respectively. We have seen that Georgian lexicalizes number concepts in (1b), while Fulfulde lexicalizes those in (1c). In other words, the lexicon size of Georgian is 11, and that of Fulfulde 6. The average morphosyntactic complexities of utterances of Georgian and Fulfulde (computed according to the formula in (2)) are 1.13 and 1.28 respectively. In other words, Fulfulde has a smaller lexicon size than Georgian, but Georgian has a lower average morphosyntactic complexity of numerals than Fulfulde.

Finally, it is interesting to notice another property of the data in Figure 1. Natural languages lie along the region of the Pareto frontier where both average morphosyntactic complexity of numerals and lexicon size are relatively low. In other words, it would seem that there are levels of average morphosyntactic complexity of numerals and lexicon size above which natural languages are, in a sense, not willing to go as far as recursive numeral systems are concerned, even if it is still possible to achieve Pareto-optimality with such higher levels of average morphosyntactic complexity of numerals and lexicon size. This would in turn suggest that the lexicon size/average morphosyntactic complexity of numerals trade-off optimization is not the only pressure shaping numeral systems: there seem to exist additional pressures, namely to keep lexicon size and average morphosyntactic complexity of numerals below a certain level, pushing towards certain regions of the Pareto frontier.

complexity. In other words, they would be shifted slightly down-right-wards in Figure 1. It is easy to see that they would remain close to the Pareto frontier under such a shift.

3 Discussion: What number meanings are lexicalized across languages?

In Section 2.1, we have seen that recursive numeral systems differ in terms of which number meanings they lexicalize as monomorphemic expressions, but that most languages lexicalize the following numbers from the range 1-99: (1) the first n number meanings, with n varying across languages, often the first five or the first 10, and (2) a couple of additional number meanings such as 10 and/or 20 as in (1b) or (1c). Outside of the 1-99 range, languages lexicalize number meanings such as 100, 1000, 1000000 as morphemes of category M . Why don't languages lexicalize only the first n numbers, which have the highest prior probability according to the probability distribution in (3)? According to the results of the experiment reported in Section 2, natural languages' recursive numeral systems trade off optimally between lexicon size and average morphosyntactic complexity of numerals. This suggests that the reason why languages lexicalize number meanings such as 10, 20, 100, 1000, 1000000 is because these allow to construct numerals denoting large(r) numbers using very few morphemes, in a way that optimizes the trade-off between lexicon size and average morphosyntactic complexity of numerals.

One may wonder however to what extent the reduction in average morphosyntactic complexity that lexicalizing number meanings such as 10 and/or 20 affords is significant given the probability distribution in (3), according to which large numbers rarely need to be communicated, and whether the result reported in Section 2 is solely a consequence of the fact that languages often lexicalize the first few number meanings. In other words, does lexicalizing the first n number meanings suffice for a nearly optimal lexicon size/average morphosyntactic complexity of numerals trade-off?

To investigate this, we generated a space of artificial languages with the following properties: (a) they lexicalize the first n numbers, with n between 2 and 10; (b) in addition to these, they lexicalize at most two other numbers smaller than 100; (c) at most 2 of number meanings they lexicalize are lexicalized as morphemes of category M . We sampled 50000 languages from this space⁸ and analyzed their lexicon size/average morphosyntactic complexity of numerals trade-off. We plot them in Figure 2 (*artificial_first_n* languages) against the natural and artificial languages from the main experiment. *Artificial_first_n* languages are not all clustered close to the Pareto frontier (black curve) in Figure 2, which demonstrates that lexicalizing the first n numbers does not guarantee a nearly optimal lexicon size/average morphosyntactic complexity of numerals trade-off.

This analysis shows that lexicalizing number meanings such as 10 and/or 20 in addition to the first n number meanings are good choices in terms of the optimality of the trade-off between lexicon size and average morphosyntactic complexity of numerals. Of course, these might not be the *only* good choices: what is the space of optimal choices for lexicalization of number meanings and whether additional pressures are at play pushing natural languages towards a subset of that space of optimal choices is an interesting avenue for future work.

⁸Due to computational constraints, it was not possible to analyze the entire space of such languages. Numerals for numbers 1-99 in the sampled languages are generated in the same way as for artificial languages in the main experiment (cf. Section 2.3).

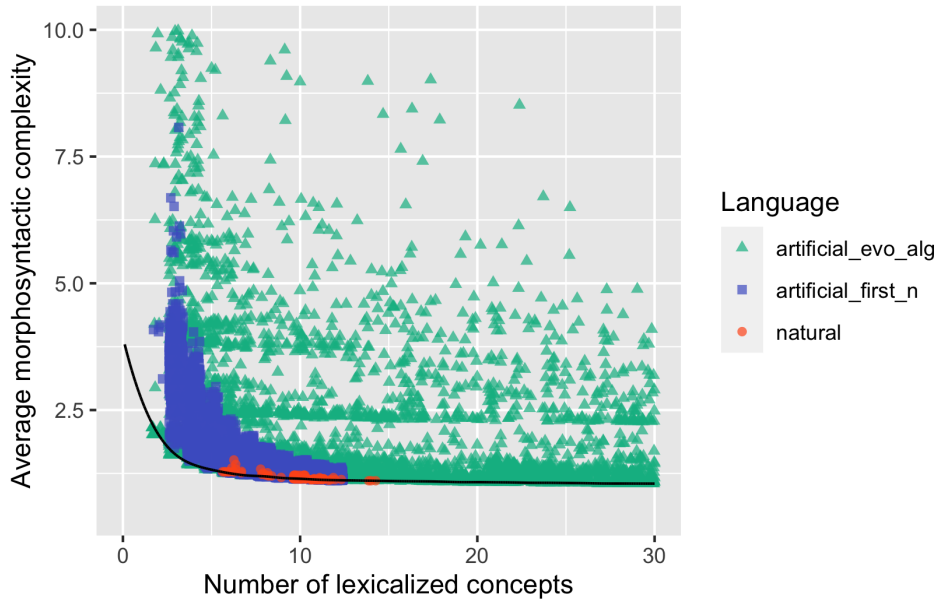


Figure 2: Lexicon size and average morphosyntactic complexity of natural languages compared to artificial languages generated via an evolutionary algorithm (*artificial_evo_lang*) as in Figure 1, with 50 000 artificial languages which lexicalize the first n numbers and at most 2 additional number meanings (*artificial_first_n*). *Artificial_first_n* languages are not all clustered close to the Pareto frontier (black curve), which demonstrates that lexicalizing the first n numbers does not guarantee a nearly optimal lexicon size/average morphosyntactic complexity trade-off.

4 Discussion: Previous work on minimization of utterance complexity

As mentioned in Introduction, several other lines of work incorporate the idea that speakers attempt to minimize complexity of their utterances. We will now discuss these proposals in detail, and explain in what way they differ from the proposal we are pursuing.

Zipf (1949) and related work Multiple studies have found that meanings that are conveyed more frequently tend to be associated to shorter forms (with length often operationalized as the number of phonemes) (Zipf 1949, Piantadosi et al. 2011, Haspelmath 2021, Mollica et al. 2021). For instance, in a recent work, Mollica et al. (2021) conduct two types of analyses: the first type of analysis shows that languages lexicalize meanings which allow them to optimize the simplicity/informativeness trade-off; the second type of analysis shows that the forms attached to lexicalized meanings are such that average phonetic length of utterances is minimized. Association of shorter forms to more frequent meanings is an instance of Zipf’s principle of least effort (Zipf 1949), according to which humans are prone to spending the least amount of effort to accomplish a task.

The pressure to minimize average morphosyntactic complexity of utterances can be viewed as another instance of the same principle. The novelty of our proposal is in showing that this pressure plays a role in determining which meanings get lexicalized, and not only how long forms are associated with lexicalized meanings. According to our proposal, this pressure on its own doesn’t suffice to explain which meanings are lexicalized across languages: we have shown that, in semantic domains in which productive morphosyn-

tax enables precise communication, languages lexicalize meanings which allow them to optimize the trade-off between this pressure and the pressure to minimize lexicon size.

Carcassi and Sbardolini (2022) Another recent proposal gives a central role to utterance complexity in explaining typological patterns, focusing on the semantic domain of Boolean connectives (Carcassi and Sbardolini 2022). They aim to explain the so-called *nand*-puzzle: across languages, connectives *and* and *or* are often lexicalized; in addition to them, the negated disjunction *nor* is sometimes lexicalized, but a negated conjunction **nand* is never lexicalized (Horn 1972). Carcassi and Sbardolini (2022) propose that the systems of Boolean connectives are optimal solutions to two pressures. The first pressure is the desire to minimize production effort when expressing observations (what we have called *average morphosyntactic complexity*). For instance, if languages don't lexicalize *nand*, they have to use syntactically more complex *not (A and B)* instead of *A nand B* to express the observation that at least one of A and B is false: lexicalizing *nand* thus reduces average morphosyntactic complexity. The second pressure, labeled *conceptual complexity* in Carcassi and Sbardolini (2022), relates to how sentences with connectives are used to update the information shared by conversational participants (= *context*), which is modeled as the set of possible worlds (Stalnaker 1978, 1999). According to their proposal, upon accepting a sentence *A nor B*, the conversational participants update the context *C* as follows: first, they restrict *C* to the set of worlds where A is false, and then they further restrict *C* to the set of worlds where B is false. On the other hand, upon accepting a sentence *A nand B*, the conversational participants update the context *C* as follows: first they consider a hypothetical context *C'*, obtained by restricting the original context *C* to the set of worlds where A is true, followed by a further restriction of *C'* to the set of worlds where B is true. They then remove the hypothetical context *C'* from the original context *C*. Carcassi and Sbardolini (2022) assume that how context is updated by *A nand B* is more conceptually complex than how context is updated by *A nor B* because of the creation of the hypothetical context *C'* in the former but not in the latter case. They propose that languages are under pressure to minimize the total conceptual complexity of contextual updates of the connectives they lexicalize. They argue that the *nand*-puzzle can be explained on the assumption that languages optimize the trade-off between average morphosyntactic complexity and the total conceptual complexity of contextual update procedures of lexicalized connectives.

Our proposal shares with Carcassi and Sbardolini (2022) the idea that languages are under the pressure to minimize average morphosyntactic complexity. Where the two proposals differ is what the competing pressure is: in our case, the pressure to minimize how many meanings are lexicalized, while in Carcassi and Sbardolini (2022) the competing pressure relates to another aspect of language use, namely, how conversational participants proceed to update contextual information when they hear an expression.

Memory/computation trade-off optimization in morphological processing The work on psycholinguistic processing of morphologically complex expressions has studied the trade-off between memory and computation (e.g. Frauenfelder and Schreuder (1992), and much related work, including more recently e.g., Kuperman, Bertram, and Baayen (2010) and O'Donnell (2015)). The central question in that line of work is: which morphologically complex linguistic expressions are retrieved from memory as a chunk, rather than computed morphosyntactically at each use? Of course, language users must store in memory morphologically simple linguistic expressions, such as 'walk', and the past tense

morpheme ‘-ed’. But do they also store in their memory ‘walked’, so that when they construct a sentence ‘John walked’, they retrieve ‘walked’ as a chunk, rather than construct it as ‘walk’+‘ed’? One response is that complex expressions are stored in memory in a way that optimizes the memory/computation trade-off in language processing. Simplifying a lot, if some morphologically complex expression needs to be used very frequently, we are more likely to store it in memory as a chunk and retrieve it as a whole in language use. At least some English speakers (those who use the linguistic expression ‘walked’ a lot) will end up storing it as a chunk in their memory.

That line of work has thus established the existence of memory/computation trade-off optimization bias in morphological processing. This connects to our work in an interesting way. What we have shown is that lexicon size/average morphosyntactic complexity of utterances trade-off optimization — which is essentially memory/computation trade-off optimization — ends up shaping the primitives (monomorphemic expressions) of the linguistic system, and explains typological diversity in what these primitives are. It is conceivable that this is an outcome of the accumulated effect in the course of language evolution of the aforementioned morphological processing bias.

5 Discussion: Xu et al. (2020) and three pressures

Xu et al. (2020) analyze the simplicity/informativeness trade-off in 24 *restricted* and 6 *recursive* numeral systems. Restricted numeral systems don’t have numerals for all numbers: most of them have numerals for only the first few numbers, and use a quantifier such as *many* for any higher number. For instance, the language Krenak only has numerals for numbers 1-3 (Hammarström 2010), and the language Rama only has numerals for numbers 1-5 (Grinevald 1990). Furthermore, the few numerals in restricted numeral systems are often monomorphemic. Recursive numeral systems are considered by Xu et al. (2020) to be maximally informative when it comes to communicating about number meanings, while restricted numeral systems have lower degrees of informativeness. On the other hand, according to Xu et al.’s (2020) approach to measuring complexity, the six studied recursive numeral systems are more complex than most of the 24 studied restricted numeral systems. Simplifying somewhat, this is because they assume that the complexity measure of a language should incorporate both the complexity of the lexicon and the complexity of morphosyntactic rules. As recursive numeral systems always have morphosyntactic rules for building numerals but restricted numeral systems often don’t have any, recursive numeral systems tend to have a greater measure of complexity than restricted numeral systems. Recursive numeral systems are thus more complex and more informative than restricted numeral systems in Xu et al.’s (2020) study. Xu et al. (2020) further argue that natural languages’ numeral systems optimize the simplicity/informativeness trade-off, without making a distinction between restricted and recursive languages.

Their conclusion may seem to be in tension with our findings. However, a careful examination of Xu et al.’s (2020) results reveals that, while restricted numeral systems are indeed close to being Pareto-optimal in trading off simplicity and informativeness, recursive numeral systems are not. Given that recursive numeral systems are maximally informative, if they were optimizing the simplicity/informativeness trade-off, we would expect to find them in the proximity of the minimally complex numeral system which is maximally informative. According to Xu et al.’s (2020) results (cf. Figure 4b in Xu




Figure 3: Reproduction of Figure 4b from Xu et al. (2020) plotting complexity and communicative cost (the opposite of simplicity and informativeness) measures of natural languages’ numeral systems (red, green and blue circles) and artificial (hypothetical) numeral systems. Blue circles represent the 6 natural languages’ recursive numeral systems, and the blue line are hypothetical recursive numeral systems. The Pareto-optimal recursive numeral system is the left-most point of the blue line, and natural languages’ recursive numeral systems are not clustered close to it.

et al. 2020, reproduced here as Figure 3), the maximally informative system with minimal complexity has complexity ≈ 50 . Strikingly, the six recursive numeral systems examined by Xu et al. (2020) have much higher complexity than that, ranging between ≈ 75 and ≈ 150 . The results of Xu et al. (2020) thus also do not support the hypothesis that recursive numeral systems optimize the simplicity/informativeness trade-off — in other words, there is no tension between the results of Xu et al. (2020) and our findings. Importantly, however, Xu et al.’s (2020) conclusions still hold for restricted numeral systems which are much closer to being Pareto-optimal than recursive numeral systems (cf. Figure 3). We will now discuss the implications of their result for restricted systems and our result for recursive systems for the larger question: *what pressures determine which meanings get lexicalized across semantic domains and languages?*

Recursive and restricted numeral systems are, in a sense, two extremes. In recursive numeral systems, languages allow to unambiguously express all meanings if we consider both morphosyntactically simple and morphosyntactically complex expressions. Such languages lexicalize those meanings which allow them to optimize the trade-off between lexicon size and average morphosyntactic complexity of numerals.

In restricted numeral systems, number meanings are typically conveyed using one of the few monomorphemic numerals, and in most cases it is not possible to unambiguously single out a number meaning with an expression of a language. Such languages lexicalize those meanings which allow them to optimize the simplicity/informativeness trade-off.

However, many semantic domains in many languages are arguably in-between these two extremes. Consider for example colors: morphosyntactically complex expressions can improve how precisely we communicate about colors compared to the communication with monomorphemic color terms only, but the improvement is arguably limited (*e.g.*, *blue-green*). What approach should one pursue to explain which meanings get lexicalized in such domains?

The results on restricted and recursive numeral systems provide evidence that (at least) three pressures are shaping which meanings get lexicalized across languages: minimize complexity of the lexicon, minimize average morphosyntactic complexity of utterances and maximize informativeness. It is plausible to expect that these pressures are not specific to the semantic domain of number, i.e., that they are applicable across semantic domains and languages. We thus propose the following hypothesis to be pursued in future work: *languages lexicalize those meanings which allow them to be optimal solutions to the three pressures*. In the extreme case where any meaning can be conveyed precisely and informativeness is maximal (as in recursive numeral systems), this reduces to finding optimal solutions to the lexicon size/average morphosyntactic complexity of utterances trade-off problem, for which we have seen evidence in the present study. In the other extreme case where interlocutors communicate about meanings from a semantic domain using single monomorphemic words only (as is arguably most often the case in restricted numeral systems), this reduces to finding optimal solutions to the simplicity of the lexicon/informativeness trade-off problem, as in [Xu et al. \(2020\)](#).

6 Conclusion

In this paper, we ask what explains which meanings get lexicalized across languages. We pursue the explanation according to which languages lexicalize meanings which allow them to support efficient communication. We however argue for a refinement of what it means for a language to support efficient communication. In particular, the standard approach to communicative efficiency — the simplicity/informativeness trade-off optimization approach — cannot explain which meanings get lexicalized in semantic domains in which lexicon size and informativeness are not in direct competition. Using recursive numeral systems as a case study, we have argued that in such domains languages optimize the lexicon size/average morphosyntactic complexity of utterances trade-off.

Our work in combination with previous work thus evidences that there are (at least) three pressures shaping lexicons of natural languages: minimizing lexicon size, maximizing informativeness, and minimizing average morphosyntactic complexity of utterances. A more general proposal for how supporting efficient communication shapes which meanings get lexicalized across languages may thus be that languages lexicalize those meanings which allow them to solve the trade-off problem between these three pressures in a (nearly) optimal way.

Acknowledgements

For helpful discussions, we wish to thank Emmanuel Chemla and Roni Katzir, as well as the audiences of the CoSaQ closing workshop in Amsterdam, LINGUAE seminar at École Normale Supérieure in Paris, Tel Aviv University Linguistics Colloquium and the 23rd Amsterdam Colloquium. This research was supported by a postdoctoral fellowship

awarded to MD by the Azrieli Foundation, and by the funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. STG 716230 CoSaQ.

Additional information

A reduced version of the present work was published as [Denić and Szymanik \(2022\)](#).

References

- Jon Barwise and Robin Cooper. Generalized quantifiers and natural language. In *Philosophy, language, and artificial intelligence*, pages 241–301. Springer, 1981.
- Brent Berlin and Paul Kay. *Basic color terms: Their universality and evolution*. Univ of California Press, 1969.
- Lisa Bylina and Rick Nouwen. Numeral semantics. *Language and Linguistics Compass*, 14(8):e12390, 2020.
- Fausto Carcassi and Giorgio Sbardolini. Assertion, denial, and the evolution of boolean operators. *Mind & Language*, 2022.
- Emmanuel Chemla, Brian Buccola, and Isabelle Dautriche. Connecting content and logical words. *Journal of Semantics*, 36(3):531–547, 2019.
- Bernard Comrie. Numeral bases. In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig, 2013. URL <https://wals.info/chapter/131>.
- Bernard Comrie, Matthew S. Dryer, David Gil, and Martin Haspelmath. Introduction. In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig, 2013. URL <https://wals.info/chapter/s1>.
- Stanislas Dehaene and Jacques Mehler. Cross-linguistic regularities in the frequency of number words. *Cognition*, 43(1):1–29, 1992.
- Milica Denić and Jakub Szymanik. Recursive numeral systems optimize the trade-off between lexicon size and average morphosyntactic complexity. In *Proceedings of the 23rd Amsterdam Colloquium*, 2022.
- Milica Denić, Shane Steinert-Threlkeld, and Jakub Szymanik. Complexity/informativeness trade-off in the domain of indefinite pronouns. In *Semantics and linguistic theory*, volume 30, pages 166–184, 2021.
- Milica Denić, Shane Steinert-Threlkeld, and Jakub Szymanik. Indefinite pronouns optimize the simplicity/informativeness trade-off. *Cognitive Science*, 2022.
- Danny Fox. Antecedent-contained deletion and the copy theory of movement. *Linguistic inquiry*, 33(1):63–96, 2002.
- Uli H Frauenfelder and Robert Schreuder. Constraining psycholinguistic models of morphological processing and representation: The role of productivity. In *Yearbook of morphology 1991*, pages 165–183. Springer, 1992.
- Peter Gärdenfors. *The Geometry of Meaning*. The MIT Press, 2014.
- Colette G Grinevald. A grammar of rama. 1990.
- Martin Hackl. On the grammar and processing of proportional quantifiers: most versus more than half. *Natural Language Semantics*, 17(1):63–98, 2009.
- Harald Hammarström. *Rarities in numeral systems*. De Gruyter Mouton, 2010.

- Martin Haspelmath. Explaining grammatical coding asymmetries: Form–frequency correspondences and predictability. *Journal of Linguistics*, 57(3):605–633, 2021.
- Laurence Robert Horn. *On the semantic properties of logical operators in English*. University of California, Los Angeles, 1972.
- Tim Hunter and Jeffrey Lidz. Conservativity and learnability of determiners. *Journal of Semantics*, 30(3):315–334, 2013.
- James R Hurford. *The linguistic theory of numerals*, volume 16. Cambridge University Press, 1975.
- James R Hurford. A performed practice explains a linguistic universal: Counting gives the packing strategy. *Lingua*, 117(5):773–783, 2007.
- Gerhard Jäger. Natural Color Categories Are Convex Sets. In Maria Aloni, Harald Bastiaanse, Tikitou de Jager, and Katrin Schulz, editors, *Logic, Language, and Meaning: Amsterdam Colloquium 2009*, pages 11–20. 2010.
- Edward L Keenan and Jonathan Stavi. A semantic characterization of natural language determiners. *Linguistics and philosophy*, pages 253–326, 1986.
- Charles Kemp and Terry Regier. Kinship categories across languages reflect general communicative principles. *Science*, 336(6084):1049–1054, 2012. doi: 10.1126/science.1218811.
- Charles Kemp, Yang Xu, and Terry Regier. Semantic typology and efficient communication. *Annual Review of Linguistics*, 4:109–128, 2018. doi: 10.1146/annurev-linguistics-011817-045406.
- Victor Kuperman, Raymond Bertram, and R Harald Baayen. Processing trade-offs in the reading of dutch derived words. *Journal of Memory and Language*, 62(2):83–97, 2010.
- Mora Maldonado and Jennifer Culbertson. Person of interest: Experimental investigations into the learnability of person systems. *Linguistic Inquiry*, 53(2):295–336, 2022.
- Francis Mollica, Geoff Bacon, Noga Zaslavsky, Yang Xu, Terry Regier, and Charles Kemp. The forms and meanings of grammatical markers support efficient communication. *Proceedings of the National Academy of Sciences*, 118(49), 2021.
- Timothy J O’Donnell. *Productivity and reuse in language: A theory of linguistic computation and storage*. MIT Press, 2015.
- Stanley Peters and Dag Westerståhl. *Quantifiers in language and logic*. Oxford University Press, 2006.
- Steven T Piantadosi. A rational analysis of the approximate number system. *Psychonomic bulletin & review*, 23(3):877–886, 2016.
- Steven T Piantadosi, Harry Tily, and Edward Gibson. Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Sciences*, 108(9):3526–3529, 2011.
- Terry Regier, Charles Kemp, and Paul Kay. Word meanings across languages support efficient communication. In Brian MacWhinney and William O’Grady, editors, *The Handbook of Language Emergence*, volume 87, pages 237–263. Wiley-Blackwell, Hoboken, NJ, 2015. doi: 10.1002/9781118346136.ch11.
- Jacopo Romoli. A structural account of conservativity. *Semantics-Syntax Interface*, 2(1):28–57, 2015.
- Benjamin Spector. Bare numerals and scalar implicatures. *Language and Linguistics Compass*, 7(5):273–294, 2013.
- Robert C Stalnaker. Assertion. In *Pragmatics*, pages 315–332. Brill, 1978.
- Robert C Stalnaker. *Context and content: Essays on intentionality in speech and thought*.

- Clarendon Press, 1999.
- Shane Steinert-Threlkeld. Quantifiers in natural language optimize the simplicity/informativeness trade-off. In *Amsterdam Colloquium 2019*, 2019. URL http://events.illc.uva.nl/AC/AC2019/uploaded_files/inlineitem/Steinert-Threlkeld_Quantifiers_in_natural_language_.pdf.
- Shane Steinert-Threlkeld. Quantifiers in natural language: Efficient communication and degrees of semantic universals. *Entropy*, 23(10):1335, 2021.
- Shane Steinert-Threlkeld and Jakub Szymanik. Learnability and semantic universals. *Semantics and Pragmatics*, 2018.
- Shane Steinert-Threlkeld and Jakub Szymanik. Ease of learning explains semantic universals. *Cognition*, 195:104076, 2020.
- Wataru Uegaki. The informativeness/complexity trade-off in the domain of boolean connectives. *Linguistic Inquiry*, pages 1–39, 2022.
- Yang Xu, Terry Regier, and Barbara C. Malt. Historical semantic chaining and efficient communication: The case of container names. *Cognitive Science*, 40(8):2081–2094, 2016. doi: 10.1111/cogs.12312.
- Yang Xu, Emmy Liu, and Terry Regier. Numeral systems across languages support efficient communication: From approximate numerosity to recursion. *Open Mind*, 4: 57–70, 2020. doi: 10.1162/opmi_a.00034.
- Noga Zaslavsky, Charles Kemp, Terry Regier, and Naftali Tishby. Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, 115(31):7937–7942, 2018.
- Noga Zaslavsky, Mora Maldonado, and Jennifer Culbertson. Let’s talk (efficiently) about us: Person systems achieve near-optimal compression. In *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*, 2021.
- George Kingsley Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, Cambridge, MA, 1949.