

Complexity and learnability in the explanation of semantic universals of quantifiers

Iris van de Pol, Shane Steinert-Threlkeld, Jakub Szymanik

{I.P.A.vandePol, S.N.M.Steinert-Threlkeld, J.K.Szymanik}@uva.nl

Institute for Logic, Language and Computation, University of Amsterdam

Abstract

Despite wide variation among natural languages, there are linguistic properties universal to all (or nearly all) languages. An important challenge is to explain why these linguistic universals hold. One explanation employs a learnability argument: semantic universals hold because expressions that satisfy them are easier to learn than those that do not. In an exploratory study we investigate the relation between learnability and complexity and whether the presence of semantic universals for quantifiers can also be explained by differences in *complexity*. We develop a novel application of (approximate) Kolmogorov complexity to measure fine-grained distinctions in complexity between different quantifiers. Our results indicate that the monotonicity universal can be explained by complexity while the conservativity universal cannot. For quantity we did not find a robust result. We also found that learnability and complexity pattern together in the monotonicity and conservativity cases that we consider, while that pattern is less robust in the quantity cases.

Keywords: semantic universals; generalized quantifiers; Kolmogorov complexity; learnability

Introduction

Even though there is huge variability between natural languages, they still share many common features. Such universal linguistic properties have been found at many levels of analysis: phonology (Hyman, 2008), syntax (Chomsky, 1965; Newmeyer, 2008), and semantics (Barwise & Cooper, 1981). Confronted with attested linguistic universals, the question naturally arises: why these properties? What explains the presence of the particular observed universals across languages?

In search of an explanation in terms of the interaction between linguistics and the specifics of human cognition, several theories have presented some form of learnability as an explanation of the presence of semantic universals (see, e.g., Barwise & Cooper, 1981; Keenan & Stavi, 1986; Szabolcsi, 2010). Recently, Steinert-Threlkeld and Szymanik (in press, henceforth ST&S) provided evidence for a version of this learnability hypothesis by using recurrent neural networks as a model for learning and applying this to several different semantic universals.

In this paper, we ask whether these semantic universals could also be explained by some measure of *complexity*, and whether this provides similar results as using a measure of learnability. It is a common expectation that there will be a connection between learnability and complexity and many theories of learning are built around such a connection (Tiede, 1999; Hsu, Chater, & Vitányi, 2013). At the same time, there

are few examples that provide evidence for this expectation in concrete cognitive tasks. In particular, it remains open whether a connection between learnability and complexity exists for independently motivated measures of each of these factors in specific domains. In the present work, we study the meaning of generalized quantifiers and compare their complexity (in a sense to be made precise) with the learnability results of ST&S.

The complexity of generalized quantifiers has been intensively studied using methods from logic, automata theory, and computational complexity.¹ However, as we will explain in more detail in a later section, none of these theories have developed a notion of complexity that applies to all quantifiers and can capture the difference between those that are attested and non-attested in natural language. To overcome these limitations, in this paper we propose to evaluate the complexity of quantifiers from an information-theoretic perspective. This perspective has already proven fruitful as an explanatory device in linguistics (Gibson et al., 2019). More specifically, we suggest to adopt (approximate) Kolmogorov complexity (Li & Vitányi, 2008) as a measure for the complexity of quantifiers. Kolmogorov complexity roughly measures how much regularity exists in a string, which enables it to be described by a shorter program that generates it. It is not implausible that universals will have the function of creating “patterns” that enable such compression.

The paper is structured as follows. In the next section, we present generalized quantifier theory and the semantic universals that we will discuss. We also discuss a recent explanation of semantic universals in terms of learnability and previous approaches to measuring the complexity of quantifiers and their limitations with respect to the current study. Following that, we introduce Kolmogorov complexity and a tractable approximation to it, and we explain how we apply this measure to binary encodings of quantifiers. In the section after that, we apply this complexity measure to the same pairs of quantifiers as in the recent learnability study to see (i) whether—in addition to learnability—some of the attested semantic universals can be explained by differences in complexity and (ii) whether complexity and learnability pattern together. We conclude by discussing the results and outlining future work.

¹See Szymanik (2016) for an overview.




Figure 1: An example of a quantifier model $\mathcal{M} = \langle M, A, B \rangle$ with 10 objects, shown as a vendiagram. This model verifies quantifiers *some* and *most*, but does not verify *all*

Quantifiers and their universal properties

Quantifiers are the semantic objects that are expressed by determiners, such as *some*, *most*, or *all*. Determiners are expressions that can combine with common nouns and a verb phrase in simple sentences of the form *Det N VP*, like “some houses are blue”. We assume a distinction of the determiners into the grammatically simple (e.g. *some*, *few*, *many*) and the grammatically complex (e.g. *at least 6* or *at most 2*, an even number of).

We use the framework of generalized quantifiers to represent the meaning of quantifiers as sets of sets. In particular, determiners denote type $\langle 1, 1 \rangle$ generalized quantifiers, which are sets of models of the form $\mathcal{M} = \langle M, A, B \rangle$, where M is the domain of the model, and A, B are two unary predicates (that is: $A, B \subseteq M$).² See Figure 1 for an illustration. This is an extensional representation of meaning, in which a quantifier is defined as the class of all models satisfying a given property (corresponding to the situations in which a simple sentence with that quantifier would be true). For a given model, \mathcal{M} , and quantifier Q we write $Q \in \mathcal{M}$ if and only if: $\mathcal{M} \models Q(A, B)$. For example, the meaning of the quantifiers *some*, *most*, and *every* can then be represented as follows:

$$\begin{aligned} \llbracket \text{some} \rrbracket &= \{ \langle M, A, B \rangle : |A \cap B| \neq \emptyset \}, \\ \llbracket \text{most} \rrbracket &= \{ \langle M, A, B \rangle : |A \cap B| > |A \setminus B| \}, \\ \llbracket \text{every} \rrbracket &= \{ \langle M, A, B \rangle : A \subseteq B \}. \end{aligned}$$

The semantic universals that we consider build on specific properties of generalized quantifiers, namely *monotonicity*, *quantity*, and *conservativity*. Let Q be a generalized quantifier. Then we call Q *monotone* if it is either upward or downward monotone, which is defined as follows. Q is *upward monotone* := if $\langle M, A, B \rangle \in Q$ and $B \subseteq B'$, then $\langle M, A, B' \rangle \in Q$. Q is *downward monotone* := if $\langle M, A, B \rangle \in Q$ and $B \supseteq B'$, then $\langle M, A, B' \rangle \in Q$. Barwise and Cooper (1981) formulate and defend the following semantic universal:

MONOTONICITY UNIVERSAL: All simple determiners are monotone.

²For a textbook treatment of generalized quantifiers see Peters and Westerståhl (2006).

The property of *quantity* intuitively expresses that the meaning of a determiner only depends on the sizes; i.e., the quantity, of the relevant sets and not on the way those sets are presented or on the particular identity of the objects in those sets. Q is *quantitative* := if $\langle M, A, B \rangle \in Q$, and $A \cap B, A \setminus B, B \setminus A$, and $M \setminus (A \cup B)$ have the same cardinality (size) as their primed-counterparts, then $\langle M', A', B' \rangle \in Q$. Keenan and Stavi (1986) formulate and defend the following semantic universal³:

QUANTITY UNIVERSAL: All simple determiners are quantitative.

The property of *conservativity* intuitively expresses that a noun phrase of the form *Det N VP* is genuinely about the N and not about the VP . That is, to verify a quantifier in a quantifier model only the A 's that are B 's are relevant, not the B 's that are not A 's. Q is *conservative* := $\langle M, A, B \rangle \in Q$ if and only if $\langle M, A, A \cap B \rangle \in Q$. Barwise and Cooper (1981) formulate and defend the following semantic universal:

CONSERVATIVITY UNIVERSAL: All simple determiners are conservative.

Explaining semantic universals via learnability

The question naturally arises: can a unified explanation be given for these universals? ST&S develop the following *learnability hypothesis*: expressions satisfying semantic universals are easier to learn than those that do not.⁴ To anthropomorphize: as languages are developing, they choose to attach lexical items to easy-to-learn meanings, and rely on complex grammatical constructions and compositional interpretation thereof to express hard-to-learn meanings.

The hypothesis immediately raises a challenge: to provide a model of learning on which it's true. ST&S train recurrent neural networks to learn minimal pairs of quantifiers, one satisfying the universal and one that does not.

Figure 2 shows an example learnability result from ST&S: an upward monotone quantifier (in blue: *at least 4*) was robustly easier to learn for a neural network than a non-monotone quantifier (in red: *at least 6* or *at most 2*). Similar patterns were observed for downward monotone and quantitative quantifiers, while conservative ones were found to be no easier to learn than non-conservative ones (but were argued to arise from a different source than learnability).

These computational results provide strong support for the learnability hypothesis. The approach has also worked well in explaining universals in disparate linguistic domains: color terms (Steinert-Threlkeld & Szymanik, 2019) and responsive predicates (Steinert-Threlkeld, in press).

Previous approaches to the complexity of quantifiers

In the literature on generalized quantifiers one can find several approaches to measuring complexity. Although these measures can capture some of the cognitive difficulty of quantifier

³See also Peters and Westerståhl (2006), van Benthem (1984), and ST&S.

⁴Hints of this hypothesis may be found in (van Benthem, 1987; Peters & Westerståhl, 2006; Magri, 2015).




Figure 2: Learning curves on a neural network for the monotone *at least 4* (blue) versus *at least 6 or at most 2* (red). This was Figure 4 in ST&S.

processing,⁵ we will see that they are not fine-grained enough to explain the presence of semantic universals.

The earliest approach uses logic, analyzing which quantifiers are definable in progressively more expressive logics. Many natural language quantifiers can be expressed in elementary (i.e. first-order) logic, e.g. *some* or *at least 4*. The seminal result here is that proportional quantifiers cannot be defined in elementary logic: one needs a stronger logical system, like second-order logic, to uniformly express the meaning of, e.g., *most*.⁶ This definability criterion cannot, however, distinguish between the complexity of the quantifiers satisfying and not satisfying the universals we study. For example, *all* and *only* can be defined with elementary formulas of exactly the same form (and therefore the same complexity):

$$\begin{aligned} \text{All}(A, B) &:= \forall x(A(x) \implies B(x)) \\ \text{Only}(A, B) &:= \forall x(B(x) \implies A(x)) \end{aligned}$$

Also, both monotone and non-monotone quantifiers can be defined by formulas of the same complexity.

Johan van Benthem (1984) has proposed to study minimal computational devices (automata) corresponding to generalized quantifiers. Under this approach, some quantifiers can be associated with canonical minimal finite automata. One can then use the size of such an automaton (i.e., the number of states) as a measure of quantifier complexity. For example, the automaton for *all* has two states while the automaton for *at least 3* has four states. Other quantifiers—for example, proportional quantifiers—must be associated with more complex computational devices, like push-down automata. This measure of complexity can explain some variance in the cognitive difficulty of quantified sentence verification against pictures (Szymanik & Zajenkowski, 2010). It is, however, not suitable for our purposes. One can easily construct a minimal quantifier pair that cannot be distinguished by this complexity measure.

⁵See Szymanik (2016) for an overview.

⁶See Peters and Westerståhl (2006) for an overview.

For instance, both *all* and *only* have minimal automata with two states. One can also easily construct a family of quantifiers with the same automaton complexity containing both quantifiers satisfying and not satisfying quantity (*at least 4*, *first 3*) and monotone and non-monotone quantifiers (*at least 4*, *at least 3* or *at most 2*). An extra problem for this approach is that for push-down automata corresponding to proportional quantifiers, there is no accepted complexity measure because they do not have a definition of a minimal automaton. So the measure does not apply to all quantifiers, including ones expressed in natural language.⁷

Another well-studied approach to identify the complexity of generalized quantifiers uses the toolbox of computational complexity theory (Szymanik, 2016). It measures quantifier complexity in terms of the asymptotic growth of the computational resources needed to recognize their meaning. The problem is that computational complexity distinctions are even more crude than the previously described alternatives. Even though computational complexity distinctions have been used to theoretically delimit the borders of natural language expressivity (Ristad, 1993; Kontinen & Szymanik, 2008), these borders include both quantifiers satisfying and not satisfying the semantic universals that we are interested in.

Kolmogorov complexity of quantifiers

To investigate whether the aforementioned semantic universals can be explained by differences in complexity, we need a measure of complexity that is suited for that task. As discussed in the previous subsection, setting up the right framework for this is a non-trivial challenge, as many well-know complexity measures are limited in their ability to distinguish between quantifiers with and without the universal properties under consideration.

Therefore, in this study, we use (approximate) Kolmogorov complexity—a finer-grained measure that has not yet been explored in this domain, and we investigate its potential to explain semantic universals. Because Kolmogorov complexity is more fine-grained than the previously discussed complexity measures, it has greater promise in capturing differences in complexity between the quantifiers that we consider. It makes intuitive sense that humans would be sensitive to Kolmogorov complexity, because it is a mathematical operationalization of the notion of compressibility and various aspects of cognition can plausibly be understood in terms of data compression: storing data compactly in a way that it can be (partially) recovered. Kolmogorov complexity has been shown useful in modelling a cognitive bias towards simplicity in a variety of cognitive domains (see Chater & Vitányi, 2003; Feldman, 2016).

Kolmogorov complexity (K) measures how much an individual sequence of symbols can be compressed. When a sequence contains regularities, these regularities can be exploited to produce a shorter description of that sequence. $K(x)$

⁷This approach has also inspired learnability models (Gierasimczuk, 2005; Clark, 2010).




Figure 3: Encoding of `some` over models of size 1.

of a sequence x is defined as the length of the shortest program p that outputs x (see Li & Vitányi, 2008).⁸

The drawback is that K has been formally proven to be uncomputable. This means that there exists no algorithm that outputs $K(x)$, given x (Li & Vitányi, 2008). For this reason, we use a well-established and tractable approximation to K , that is based on the Lempel-Ziv algorithm for lossless data compression (Lempel & Ziv, 1976). The Lempel-Ziv algorithm parses a sequence x from left to right, and cuts up the sequence into subsequences. At each point it chooses the longest possible subsequence that is identical to an earlier part of the sequence, thereby identifying the number of unique subpatterns in x . The Lempel-Ziv complexity $LZ(x)$ is the number of these unique subpatterns of x . For approximate Kolmogorov complexity \tilde{K} , we use $C_{LZ}(x)$, which is defined as $\log_2(\text{len}(x)) \cdot LZ(x)$.⁹ Ziv and Lempel (1978) show that $C_{LZ}(x)$ approximates $K(x)$ in the limit; i.e. when $\text{len}(x)$ approaches infinity. Vitányi (2013) shows that, in practice, lossless compression methods give adequate results also for finite sequences. Furthermore, C_{LZ} is considered particularly adequate as a measure for \tilde{K} for shorter strings (Lesne, Blanc, & Pezard, 2009).¹⁰

To determine the approximate Kolmogorov complexity \tilde{K} of a quantifier we need to represent it as a sequence of symbols. We encode a quantifier as a binary sequence, representing the quantifier as a distribution of truth values over all models (up to a certain size). First, we enumerate all possible models. Then, given such an enumeration, we represent a quantifier by placing a 1 in the sequence for every model that verifies the quantifier and placing a 0 for every model that does not verify the quantifier. See Figure 3 for an example. Given a sequence of models, this gives a unique binary representation for every possible quantifier. Then, for a given sequence of models up to a certain maximum model size, we can determine the complexity of a quantifier Q by computing $\tilde{K}(x_Q)$ over the binary representation x_Q of Q .

⁸Formally, K is defined given a particular universal Turing machine (UTM), but, by the Invariance Theorem, K given UTM V or given UTM W will not differ more than some constant c .

⁹In particular, we use the same version of C_{LZ} as used by Dingle, Camargo, and Louis (2018), which uses the average between $LZ(x)$ and $LZ(\text{reverse}(x))$ to obtain an even more fine-grained complexity measure.

¹⁰There are also other popular lossless compression methods that can be used as approximations to K , such as `gzip` (based on LZ compression), and `bzip2` (a block-sorting compressor). Graphs comparing the LZ and `gzip2` complexity of the quantifiers that we considered can be found at <https://tinyurl.com/quantifierLZ>.

This framework allows us to compare the complexity of different quantifiers and investigate whether semantic universals might be explained by differences in complexity. In doing this, we are not interested in the absolute complexity values of the quantifiers but in the difference in complexity between a quantifier that satisfies a universal and its minimally differing counterpart that does not satisfy that universal. To make any such comparison across quantifiers, we need to fix an enumeration over quantifier models and use that as the base for our quantifier representations.

One way of doing that would be to take a random enumeration over quantifier models. Unfortunately, for our purpose, this is not a suitable method. For a random sequence, the complexity of a quantifier is mainly determined by the uniformity of that quantifier (defined by the ratio of 1's versus 0's in the quantifier representation).¹¹ When the uniformity of a quantifier is the main determiner for its complexity, differences between the complexity of two quantifiers might not reflect differences due to the presence or absence of a particular universal property.

For our purpose, choosing a structured sequence over models is more suitable than taking a random sequence. The intuition behind this can be understood as follows. If a quantifier that satisfies a universal has lower \tilde{K} complexity than its minimally differing counterpart, then this will be because the universal property causes a regularity in the distribution of truth values across quantifier models. This difference in regularity between quantifiers could disappear when evaluating quantifiers over a random sequence of models, but it might be visible when evaluating those quantifiers over a structured and well-behaved sequence. For this reason, we evaluate our quantifiers over the lexicographic sequence of models, which is standardly used in the literature on generalized quantifiers. For robustness, we look at all 12 uniquely different possible lexicographical orderings, arising from the different ways of ordering the symbols for the four sets $A \cap B$, $A \setminus B$, $B \setminus A$, and $M \setminus (A \cup B)$.¹²

Results

With this framework in place we can now turn to our main question. To test whether approximate Kolmogorov complexity can explain the three proposed semantic universals, we

¹¹This can be understood from the fact that among all different strings of a given uniformity there are only few strings of low complexity. This is because when a string x of length n has a low complexity, this means that x can be compressed to a shorter string x' of length $n' < n$, and there are only few strings of length n' compared to the amount of strings of length n . Therefore, when taking two quantifier representations with the same uniformity, over a random sequence of models, they are likely to both have complexity values that are close to the maximum complexity for that uniformity (which are thus similar).

¹²In fact, there are in total 24 different lexicographic enumerations over the quantifier models that we use, but only 12 of them are unique, and the other 12 are the reverse of one of those 12 unique sequences. As mentioned earlier, we use a measure that takes the average between the complexity over a sequence and the complexity of the reverse of that sequence. So this leaves 12 lexicographical sequences over which we can compute this measure.

look at minimally differing pairs of quantifiers in which one satisfies the universal and the other does not. To compare our complexity results with the learnability results of ST&S, we test the same pairs of quantifiers. Let $x_{i,Q}$ be the binary representation of quantifier Q , based on a sequence of all models up to size i . For each quantifier Q , and for each model size i from 1 to 10, we computed $C_{LZ}(x_{i,Q})$. We repeated this for all 12 lexicographical model sequences. For each pair we plotted the mean complexity against the maximum model size (with confidence intervals), and we compared the differences in complexity between the two quantifiers at each maximum model size and model sequence. The code that we used for generating these data and the data themselves can be found at <https://tinyurl.com/quantifierLZ>.

Monotonicity

To test the MONOTONICITY UNIVERSAL, we looked at two quantifier pairs, one with a downward- and one with an upward-monotone quantifier. First, we compared the downward-monotone quantifier at most 3, meaning $|A \cap B| \leq 3$, with the non-monotone quantifier at least 6 or at most 2, meaning $|A \cap B| \geq 6$ or $|A \cap B| \leq 2$. The mean complexity values over all 12 lexicographical model sequences and a 95% confidence interval are plotted in Figure 4. The descriptive statistics show that for all model sizes larger than 2, monotone at least 4 has a lower complexity than non-monotone at least 6 or at most 2 (for model size 1 and 2 the differences are 0). This holds for each of the 12 different model sequences. The 12 individual plots for this pair and all the other quantifier pairs can be found at <https://tinyurl.com/quantifierLZ>.

Second, we compared the upward-monotone quantifier at least 4, meaning $|A \cap B| \geq 4$, with the non-monotone quantifier at least 6 or at most 2, meaning $|A \cap B| \geq 6$ or $|A \cap B| \leq 2$. The mean complexity values over all 12 model sequences and a 95% confidence interval are plotted in Figure 4. Exactly like for the downward-monotone quantifiers, the descriptive statistics show that that for all model sizes larger than 2, monotone at most 3 has a lower complexity than non-monotone at least 6 or at most 2 (for model size 1 and 2 the differences are 0). Again, this holds for each of the 12 different model sequences.

These complexity results show the same patterns as the learnability results of ST&S. This supports the hypothesis that, in addition to learnability, the MONOTONICITY UNIVERSAL might be explained by differences in complexity, with monotone quantifiers being less complex than non-monotone quantifiers.

Quantity

To test the QUANTITY UNIVERSAL, we looked at two quantifier pairs with a quantitative and a non-quantitative quantifier. First, we compared the quantitative quantifier at least 3, with the non-quantitative quantifier first 3. The mean complexity values over all 12 model sequences and a 95% confidence interval are plotted in Figure 5. For model size 1,




Figure 4: Complexity values for at most 3 and at least 6 or at most 2, and for at least 4 and at least 6 or at most 2. Mean values with 95% confidence interval over all 12 lexicographic model sequences

2, and 3, the differences are 0, and for model sizes 4 to 10 the descriptive statistics show that at least 3 is less complex in 59.5% of the cases, and more complex in 33.3% of the cases.

Second, we compared the quantitative quantifier at least 3, with the non-quantitative quantifier last 3. The main complexity values over all 12 model sequences and a 95% confidence interval are plotted in Figure 5. Again, for model size 1, 2, and 3, the differences are 0, while for model sizes 4 to 10 the descriptive statistics show that at least 3 is less complex in 52.4% of the cases and more complex in 42.9% of the cases.

These complexity results do not show a robust pattern. However, they do show a tendency towards the quantitative quantifiers being less complex than the non-quantitative quantifiers. In the learnability results of ST&S, the quantitative quantifiers were significantly easier to learn than the non-quantitative ones. These findings neither confirm nor disconfirm the hypothesis that, in addition to learnability, the QUANTITY UNIVERSAL could be explained by differences in complexity.

Conservativity

To test the CONSERVATIVITY UNIVERSAL, we looked at two quantifier pairs with a conservative and a non-conservative quantifier. First, we compared the conservative quantifier most, meaning $|A \cap B| > |A \setminus B|$, with the non-conservative quantifier M, meaning $|A| > |B|$. The mean complexity values over all 12 model sequences and a 95% confidence interval are plotted in Figure 6. The descriptive statistics show that that for all model sizes and for all model sequences, conservative most has exactly the same complexity as non-conservative M.




Figure 5: Complexity values for at least 3 and first 3, and for at least 3 and last 3.

Second, we compared the conservative quantifier *not all*, meaning $A \not\subseteq B$, with the non-conservative quantifier *not only*, meaning $B \not\subseteq A$. Again, the mean complexity values over all 12 model sequences and a 95% confidence interval are plotted in Figure 6. For model size 1 to 10 descriptive statistics show that *not all* is more complex in 55.9% of the cases and less complex in 40.8% of the cases.

These results do not support the hypothesis that the CONSERVATIVITY UNIVERSAL can be explained by differences in complexity. However, these complexity results do show the same patterns as the learnability results of ST&S, as in their results the conservative quantifiers were of similar learnability as the non-conservative ones. This, however, does not constitute a counterexample to the explanation of the universals via learnability. As explained by ST&S one should not expect the difference between conservative and non-conservative quantifiers under their framework. This universal should rather be explained in terms of the syntax-semantics interface.¹³

Discussion

Let us take stock. We have applied tools from algorithmic information theory—in particular, approximate Kolmogorov complexity—to measure the complexity of quantifiers expressed in natural language. We did this in order to see whether the complexity of a quantifier can explain the presence of semantic universals for quantifiers, and whether these complexity results show the same patterns as existing learnability results. We found that monotone quantifiers are robustly less complex than non-monotone quantifiers, and that conservative

¹³See Romoli (2015). Hunter and Lidz (2013) observe a difference in children learning conservative vs. non-conservative quantifiers. This result, if replicated, could be due to a bias acquired by the children in earlier exposure to only conservative determiners.




Figure 6: Complexity values for most and M, and for not all and not only.

and non-conservative quantifiers have equal or similar complexity. For quantitative quantifiers we found a slight tendency towards being less complex, but this pattern was not robust. The results for monotonicity and conservativity agree with an existing explanation in terms of learnability due to ST&S, while the results on quantitativity hint in the same direction, but not robustly so.

The results of the exploratory study that we undertook are not decisive. Nevertheless, the results show substantial similarity between the complexity and learnability of quantifiers in the explanation of semantic universals. Our results for monotonicity show that approximate Kolmogorov complexity can indeed capture differences in complexity between quantifiers that could not be captured with the complexity measures from the previous approaches that we discussed. That neither complexity nor learnability distinguishes conservative from non-conservative quantifiers provides further evidence that conservativity has a different source than the other two universals, as suggested by ST&S.

Much work remains to be done. To corroborate our results, one would like to scale up beyond maximum model size of $n = 10$; how to make this computationally efficient is not a simple task. One would also like to expand the experiments beyond the minimal pair methodology employed here. In order to compare with existing results, it would be good to measure the complexity of many quantifiers and see which semantic properties best explain the complexities. The methods here could also be applied to semantic universals in other domains, to test the connection between complexity and learnability in a more general setting. Finally, one can also look at other measures of complexity: for instance, minimal derivation length in a Language of Thought for generating expressions for quantifier meanings (Piantadosi, Tenenbaum, & Goodman, 2012; Goodman, Tenenbaum, & Gerstenberg, 2015).

Acknowledgements

We thank the members of the CoCoLab and the CoLaLa at Stanford and Berkeley, respectively, for helpful discussion and Ronald de Haan for technical support. IvP was supported by Gravitation Grant 024.001.006 of the Language in Interaction Consortium from the Netherlands Organization for Scientific Research (NWO). Shane S-T and JS have received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement n. STG 716230 CoSaQ.

References

- Barwise, J., & Cooper, R. (1981). Generalized Quantifiers and Natural Language. *Linguistics and Philosophy*, 4(2), 159–219.
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in cognitive sciences*, 7(1), 19–22.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. MIT Press.
- Clark, R. (2010). On the learnability of quantifiers. In J. J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 909–922). Elsevier.
- Dingle, K., Camargo, C. Q., & Louis, A. A. (2018). Input–output maps are strongly biased towards simple outputs. *Nature Communications*, 9(1), 761.
- Feldman, J. (2016). The simplicity principle in perception and cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 7(5), 330–340. doi: 10.1002/wcs.1406
- Gibson, E., Futrell, R., Piantadosi, S. T., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How Efficiency Shapes Human Language. *Trends in Cognitive Sciences*.
- Gierasimczuk, N. (2005). The problem of learning the semantics of quantifiers. In *International tbilisi symposium on logic, language, and computation* (Vol. 4363, pp. 117–126).
- Goodman, N. D., Tenenbaum, J. B., & Gerstenberg, T. (2015). Concepts in a probabilistic language of thought. In Morgolis & Lawrence (Eds.), *The Conceptual Mind: New Directions in the Study of Concepts*. MIT Press.
- Hsu, A. S., Chater, N., & Vitányi, P. (2013). Language learning from positive evidence, reconsidered: A simplicity-based approach. *Topics in Cognitive Science*, 5(1), 35–55.
- Hunter, T., & Lidz, J. (2013). Conservativity and learnability of determiners. *Journal of Semantics*, 30(3), 315–334.
- Hyman, L. M. (2008). Universals in phonology. *The Linguistic Review*, 25(1-2), 83–137.
- Keenan, E. L., & Stavi, J. (1986). A Semantic Characterization of Natural Language Determiners. *Linguistics and Philosophy*, 9(3), 253–326.
- Kontinen, J., & Szymanik, J. (2008). A remark on collective quantification. *Journal of Logic, Language and Information*, 17(2), 131–140.
- Lempel, A., & Ziv, J. (1976). On the complexity of finite sequences. *IEEE Transactions on Information Theory*, 22(1), 75–81.
- Lesne, A., Blanc, J.-L., & Pezard, L. (2009). Entropy estimation of very short symbolic sequences. *Physical Review E*, 79(4), 046208:1–10.
- Li, M., & Vitányi, P. (2008). *An Introduction to Kolmogorov Complexity and Its Applications*. Springer.
- Magri, G. (2015). Universals on natural language determiners from a pac-learnability perspective. In *Proceedings of cogsci 2015*.
- Newmeyer, F. J. (2008). Universals in syntax. *The Linguistic Review*, 25(1-2), 35–82.
- Peters, S., & Westerståhl, D. (2006). *Quantifiers in Language and Logic*. Oxford: Clarendon Press.
- Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2012). *Modeling the acquisition of quantifier semantics: A case study in function word learnability*.
- Ristad, E. S. (1993). *The language complexity game*. MIT Press.
- Romoli, J. (2015). A Structural Account of Conservativity. *Semantics-Syntax Interface*, 2(1), 28–57.
- Steinert-Threlkeld, S. (in press). An explanation of the veridical uniformity universal. *Journal of Semantics*.
- Steinert-Threlkeld, S., & Szymanik, J. (2019). *Ease of Learning Explains Semantic Universals*.
- Steinert-Threlkeld, S., & Szymanik, J. (in press). Learnability and Semantic Universals. *Semantics & Pragmatics*. (<https://semanticsarchive.net/Archive/mQ2Y2Y2Z/LearnabilitySemanticUniversals.pdf>)
- Szabolcsi, A. (2010). *Quantification*. Cambridge: Cambridge University Press.
- Szymanik, J. (2016). *Quantifiers and Cognition: Logical and Computational Perspectives* (Vol. 96). Springer.
- Szymanik, J., & Zajenkowski, M. (2010). Comprehension of simple quantifiers. Empirical evaluation of a computational model. *Cognitive Science: A Multidisciplinary Journal*, 34(3), 521–532.
- Tiede, H.-J. (1999). Identifiability in the Limit of Context-Free Generalized Quantifiers. *Journal of Language and Computation*, 1(1), 93–102.
- van Benthem, J. (1984). Questions About Quantifiers. *The Journal of Symbolic Logic*, 49(2), 443–466.
- van Benthem, J. (1987). Toward a Computational Semantics. In P. Gardenfors (Ed.), *Generalized quantifiers: Linguistic and logical approaches* (pp. 31–71). Kluwer.
- Vitányi, P. M. (2013). Similarity and denoising. *Philosophical Transactions of the Royal Society A*, 371(1984).
- Ziv, J., & Lempel, A. (1978). Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory*, 24(5), 530–536.