

On the Evolution of Word Order

Idan Rejwan

Tel Aviv University and AI21 Labs
Tel Aviv, Israel
idanrejwan91@gmail.com

Avi Caciularu

Bar-Ilan University
Ramat-Gan, Israel
avi.c33@gmail.com

Abstract

Most natural languages have a predominant or fixed word order. For example in English the word order is usually Subject-Verb-Object. This work attempts to explain this phenomenon as well as other typological findings regarding word order from a functional perspective. In particular, we examine whether fixed word order provides a functional advantage, explaining why these languages are prevalent. To this end, we consider an evolutionary model of language and demonstrate, both theoretically and using genetic algorithms, that a language with a fixed word order is optimal. We also show that adding information to the sentence, such as case markers and noun-verb distinction, reduces the need for fixed word order, in accordance with the typological findings.

1 Introduction

Word order is a linguistic concept that refers to how syntactic elements are arranged in a sentence in different languages. Specifically, it deals with the relative position of the three prominent syntactic roles in a sentence: subject (*S*), verb (*V*), and object (*O*). The three syntactic roles can be ordered in six different arrangements: *SVO*, *SOV*, *VSO*, *VOS*, *OVS*, *OSV*. Apriori, there is no reason to assume that one order is superior to others. However, typological findings show that most natural languages have one preferred or fixed word order, with the vast majority of them belonging to the first two families: *SVO* or *SOV* (Comrie, 1989). Interestingly, languages that have developed spontaneously and did not contact other languages also have a distinct preference for particular word order, such as the Al-Sayyid Bedouin Sign Language (Sandler et al., 2005).

Another typological finding is related to case markers, which are morphological indicators (usually prefixes or suffixes) that indicate the syntactic

role of a word in a sentence. For example, the word for *the boy* in Arabic is *al-walad-u* if it is the subject and *al-walad-a* if it is the object. Apparently, in languages with a rich system of case markers, such as Arabic, Latin and Sanskrit, the word order is more flexible (Sapir, 1921; McFadden, 2004).

These findings raise several research questions. First, why is there a preferred or fixed order in the majority of natural languages? Second, why do most languages have *SVO* or *SOV* as their preferred order? Third, how can we explain the relationship between case markers and flexible word order?

This work addresses the aforementioned questions within an evolutionary perspective (Pinker and Bloom, 1990; Nowak et al., 2000). As in the evolution theory of the animal world, if a particular feature appears in many different species, it is assumed to give a functional advantage to the organisms carrying it, thus helping them survive natural selection. Similarly, linguistic phenomena that are common to all or most languages are assumed to provide a functional advantage and facilitate better communication.

Within this perspective, a plethora of works suggested to use Genetic Algorithms (GA), an optimization method inspired by evolutionary principles, in order to investigate typological findings (Batali, 1998; Kirby, 1999; Levin, 1995; Kirby, 2001). In this work, we follow this framework and use GA to study the phenomenon of word order, yet not investigated to the best of our knowledge. In particular, we attempt to answer whether a fixed word order gives a functional advantage, which may explain its emergence in most natural languages. We also provide a theoretical analysis to prove that the optimal grammar, with respect to our functional assumptions, is a grammar with fixed word order. Although it would be interesting to explore why a specific word order is preferred, for example *SVO* or *SOV*, this work addresses the

fundamental question about why languages have fixed word orders at all, rather than specific ones.

2 Method

Following Pinker and Bloom (1990), we view language as a form of communication designed to transmit a message between a speaker and a hearer. Thus, if we are to speak of an “optimal” language, it is the one that facilitates precise delivery of a message between the speaker and the listener.

In this view, the hearer observes only the sentence received from the speaker, which does not contain information regarding the relationship between the words and their syntactic roles.

Intuitively, word order might impact the way the hearer understands the sentence. For example, compare the English sentence “*we ship a man*” with “*we man a ship*”. In this case, as the words “*ship*” and “*man*” may function as both a noun and a verb, and due to the absence of case markers, the meaning is determined solely by the order of words in the sentence.

Under these assumptions, we use GA to simulate a language’s evolution, starting with no preferred word order, and test whether communication constraints enforce a fixation of particular word order.

2.1 Algorithm

The algorithm is initialized with a *lexicon*, a fixed list of 1000 words, where each word is a string of three letters. The words are randomly initialized and are given to all agents. We test both options of one lexicon for all words or two lexicons for nouns and verbs (see elaboration in section 4).

We initialize a population of 100 *grammars*, where each one is a probability distribution over the six possible word orders and is initialized to the uniform probability. Each grammar is represented by a *speaker* and a *hearer* trying to communicate a message over a noisy channel, and their success in communication measures the fitness of the grammar. In other words, for every grammar, the *speaker* produces a sentence based on the grammar, and the *hearer* determines the syntactic role of each word in the sentence.

The GA runs for 1000 successive generations. In each one, the speaker samples three random words from the lexicon with their syntactic roles. Then, he samples an order according to the probability distribution in the grammar, concatenates the sampled words according to it, and sends it to the

hearer. Then, a random noise flips every letter in the sentence with a probability of $p = 0.01$. Note that the noisy sentence might contain words that are out of the lexicon.

Next, the hearer receives the noisy sentence and matches the syntactic roles in two steps. First, for every word in the sentence, he removes the noise by looking for the nearest word in the lexicon, in terms of Levenshtein distance. Then, he outputs the most probable word order given the grammar and the words.

Finally, we measure the Hamming distance between the input and the output by counting the correct syntactic roles. For example, the distance between SVO and SOV is $2/3$ since *O*, and *V* are at different locations.

In each generation, 30% of the grammars with the smallest distances are replicated to the next generation, with some random mutations to the probability distributions - a Gaussian noise with variance 0.01 is added to each probability in the grammar, and then re-normalized, to sum up to 1. After many generations, the surviving grammars are expected to be the ones that allow optimal communication between the speaker and the hearer.

We use this method to examine whether adding information to the sentence, such as the distinction between nouns and verbs or explicit case markers, changes the optimal grammar.

We tested different values for the population size, selection rate, and mutation rate, but all the examined values yielded similar trends.

3 Theoretical Analysis

As mentioned, optimal communication occurs when the distance between the input and output is minimal. In this section, we prove that under this definition of optimality, and assuming the hearer always identifies the words in the lexicon correctly despite the noise (as indeed happens with a high probability), the following holds

Theorem 1. *Any optimal grammar has exactly one fixed word order.*

Proof. Let $X_1 \dots X_n$ and $Y_1 \dots Y_n$ be random variables of the occurrence of possible word orders, associated with the speaker and the hearer, respectively. We define

$$P(X_i = 1) = P(Y_i = 1) = p_i, \quad i = 1, \dots, n.$$

Let d_{ij} be the distance between X_i and Y_j , defined as the number of syntactic roles with different positions between X_i and Y_j . Then, the expected distance is given by

$$\mathbb{E}[d] = \sum_{(i,j)} d_{ij} P(X_i, Y_j),$$

where $P(X_i, Y_j)$ is the probability that the speaker chose order X_i (with probability of p_i) and the hearer chose order Y_j (with probability of p_j). Since the hearer has no additional information other than the observed sentence, he picks the most probable order from his grammar, independently of the speaker. Accordingly, $P(X_i, Y_j) = p_i p_j$, and therefore the expected value of the distance is given by

$$\mathbb{E}[d] = \sum_{(i,j)} p_i d_{ij} p_j.$$

Note that $p_i d_{ij} p_j \geq 0, \forall i, j$, since p_i, p_j are probabilities, and d_{ij} is a distance. Since $d_{ij} = 0$ only for $i = j$, to get the minimal expected distance 0 it should hold $\forall i \neq j$ either $p_i = 0$ or $p_j = 0$. Therefore, p_i can be greater than zero in exactly one entry. For other words, the only probability distributions that give an expected distance of zero have probability 1 for a particular word order X_i and zero for the rest. Accordingly, the optimal grammar must have fixed word order. \square

4 Experiments

We tested the model in four different scenarios.

1. **Base** - all words are taken from a single lexicon, without distinction between nouns and verbs and without case markers.
2. **N-V** - separate lexicons for nouns and verbs, which enables the hearer to distinguish between nouns and verbs, thus identifying the verb in the sentence.
3. **Case** - each word in the sentence carries a suffix indicating its syntactic role (S, V or O).
4. **N-V & Case** - both 2 and 3, i.e., there is a distinction between nouns and verbs, and every word in the sentence carries a case marker.

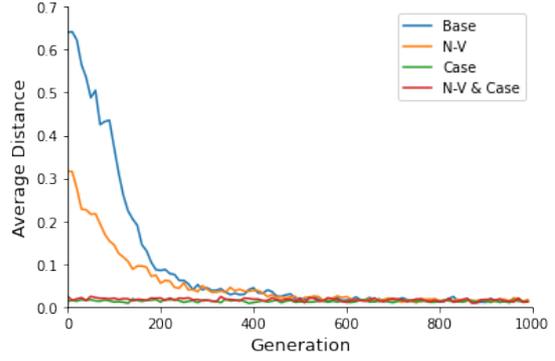


Figure 1: Average distance of sentences over generations.

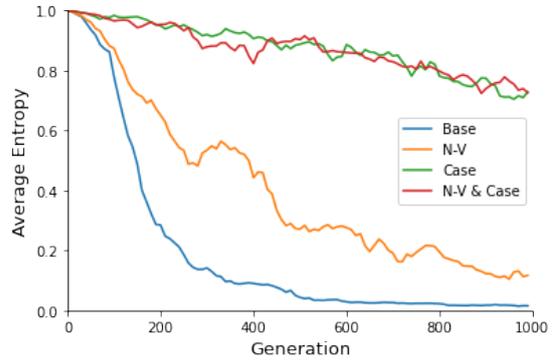


Figure 2: Average entropy of grammars over generations.

5 Results

For every scenario, we present the average distance between the input and the output among all grammars in each generation (Figure 1). It can be observed that the base scenario starts with the highest distance among the scenarios, as the sentences carry no information regarding the syntactic roles. A similar discussion could be held for the N-V scenario, in which the hearer can identify the verb in the sentence but cannot distinguish between the subject and object. In both scenarios, we observe a drop in the average distance, meaning that at some point, the hearers successfully parsed the sentences with only a small error. In the two scenarios with the case markers, we see that the distance is very close to zero even in the first generation, meaning that the hearers parsed the sentences correctly from the first generation.

To examine if the grammars converged to fixed word order, we present the grammars' average entropy in every generation (Figure 2). The entropy is defined by $H = -\sum_i p_i \log p_i$ and satisfies $0 \leq H \leq 1$ where $H = 1$ represents a uniform

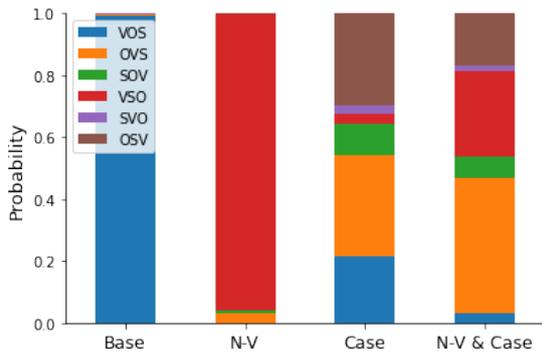


Figure 3: Word order probability after 1000 generations.

distribution (absolute randomness) and $H = 0$ represents probability 1 for one word order and zero for the rest (absolute determinism). It can be observed that the average entropy of grammars begins at 1 (absolute randomness) in all scenarios, as the probability for every word order is equal. In the base scenario, we observe a fast decline in entropy until it converges to zero, i.e., a fixed word order. A decline is also observed in the N-V scenario. However, it is slower and does not converge to zero even after 1000 generations. As for the scenarios with the case markers, in both cases, we observe that the entropy remains high even after 1000 generations, meaning that no particular word order was fixed.

We also present the grammar that had the smallest distance after 1000 generations (Figure 3). In the base and N-V scenarios, we observe that the grammar converged to an (almost) fixed word order, VOS in the base scenario, and VSO in the N-V scenario. In contrast, in both scenarios with case markers, the word order remained flexible even after 1000 generations.

6 Discussion

In the base scenario, we observed a fast fixation of single word order. Therefore, in line with the theoretical result, we conclude that given no additional information regarding the sentence’s syntactic relations, fixed word order is necessary to allow successful communication. We also observed a fixation in the N-V scenario, meaning that a distinction between nouns and verbs does not provide sufficient information for communication, and thus fixed word order is also necessary in this case.

In contrast, both scenarios with case markers did not converge to preferred word order, meaning that the case markers by themselves are enough to allow

communication. This finding is consistent with the result of artificial language learning experiments conducted on human learners (Fedzechkina et al., 2011). Note that since the case markers we used were just one letter, they could have been flipped by the random noise, leading to a wrong parsing by the hearer. However, the chances that two case markers in the same sentence will be flipped are small, and thus the hearer can parse the sentence correctly with high probability.

We note that additional experiments performed with a more significant number of syntactic roles yielded similar trends. Hence we conclude that these findings do not depend on the number of syntactic roles.

The reason why the base scenario converged to VOS and the N-V scenario converged to VSO is arbitrary, due to random drift, since there is no functional advantage to a particular word order over the others.

7 Conclusions

This work presents a simplistic functional model of language that demonstrates how a fixed word order in languages may result from a lack of information about syntactic relations.

Empirically, we used Genetic Algorithms (GA) in order to investigate whether additional information would render the fixed word order unnecessary. We showed that a distinction between nouns and verbs provides additional information to the hearer, but it is not sufficient. In contrast, case markers provide enough information to enable a flexible word order. These results are consistent with the typological findings.

Regarding the reason why most natural languages tend exhibit SVO or SOV orderings, the model used in this work does not seem to provide an explanation. It has been demonstrated that even where a single word order has been fixed, we have received various possible orders that are not necessarily one of these two. This is because we did not consider the functional advantages of certain orders over others, so the final order is entirely determined by random chance. Additional psycho-linguistic considerations are needed to explain the typological finding, as done in other studies (Gibson et al., 2013; Hengeveld et al., 2004).

References

- John Batali. 1998. Computational simulations of the emergence of grammar. *Approach to the Evolution of Language*, pages 405–426.
- Bernard Comrie. 1989. *Language universals and linguistic typology: Syntax and morphology*. University of Chicago press.
- Maryia Fedzechkina, T Florian Jaeger, and Elissa L Newport. 2011. Functional biases in language learning: Evidence from word order and case-marking interaction. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33.
- Edward Gibson, Steven T Piantadosi, Kimberly Brink, Leon Bergen, Eunice Lim, and Rebecca Saxe. 2013. A noisy-channel account of crosslinguistic word-order variation. *Psychological science*, 24(7):1079–1088.
- Kees Hengeveld, Jan Rijkhoff, and Anna Siewierska. 2004. Parts-of-speech systems and word order. *Journal of linguistics*, pages 527–570.
- Simon Kirby. 1999. Syntax out of learning: The cultural evolution of structured communication in a population of induction algorithms. In *European Conference on Artificial Life*, pages 694–703. Springer.
- Simon Kirby. 2001. Spontaneous evolution of linguistic structure—an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110.
- Michael Levin. 1995. The evolution of understanding: A genetic algorithm model of the evolution of communication. *BioSystems*, 36(3):167–178.
- Thomas McFadden. 2004. The position of morphological case in the derivation. *PhD diss., University of Pennsylvania*.
- Martin A Nowak, Joshua B Plotkin, and Vincent AA Jansen. 2000. The evolution of syntactic communication. *Nature*, 404(6777):495–498.
- Steven Pinker and Paul Bloom. 1990. Natural language and natural selection. *Behavioral and brain sciences*, 13(4):707–727.
- Wendy Sandler, Irit Meir, Carol Padden, and Mark Aronoff. 2005. The emergence of grammar: Systematic structure in a new language. *Proceedings of the National Academy of Sciences*, 102(7):2661–2665.
- Edward Sapir. 1921. *Language: An introduction to the study of speech*. Harcourt, Brace.