Reinforcement Learning Produces Efficient Case-Marking Systems

Sasha Boguraev[1] Katrin Erk[1] Kyle Mahowald[1] James Shearer[2] Stephen Wechsler[1]

[1]Department of Linguistics, The University of Texas at Austin [2]JWShearer Consulting

{sasha.boguraev,katrin.erk,kyle,wechsler}@utexas.edu
JamesWShearer@gmail.com

Abstract

Many languages mark either accusative case (for objects of transitives) or ergative case (for subjects of transitives), but some *split ergative* languages mix the two systems depending on the type of nominal. It has been noted that these languages tend towards marking the less frequent case for each nominal type. This raises the question of what mechanism could underlie the emergence of such an efficient system. We propose a model that can provide an explanation, based on a simple reinforcement learning framework and simple assumptions about asymmetries between the kinds of nominals (e.g., pronouns vs. full noun phrases) that appear in subject vs. object position.

Keywords: grammar; case-marking; communicative efficiency; language evolution

Introduction

Case systems distinguish the subject and object of a transitive clause by marking one of the nominals with a case morpheme. If only one function is marked, then which function, subject-of-transitive (A) or object (O), should it be? Which one remains unmarked, bearing the same form as the subject of intransitives (S)? Most of the world's case systems mark O with so-called accusative case, distinguishing it from unmarked A and S. Some of them mark A with so-called ergative case, distinguishing it from O and S (Fig. 1). Still others mix the two systems, depending on the type of NP. For example, in Dyirbal transitive clauses, 3rd person pronouns and lexical NPs are on an ergative system, marking subjects with the ergative case suffix *-ŋgu* (see 1a). Meanwhile 1st and 2nd person pronouns are on an accusative system, marking objects with the accusative case suffix *-na* (see 1b). The unmarked nominals have the same form as subjects of intransitives (1c).

(1) Dyirbal (Dixon, 1994, 10 and 14)

a. yabu-**ŋgu** ŋuma bura-n.
mother-**ERG** father(ABS) see-NONFUT
'Mother(A) saw father(O).'

b. ŋana n^{y}urra-**na** bura-n.
1PL(NOM) 2PL-**ACC** see-NONFUT
'We(A) saw you-all(O).'

c. ŋana/ŋuma banaga-n^{y}u.
1PL(NOM)/father(ABS) return-NONFUT
'We(S)/Father(S) returned.'

When A and S have the same form, such as *ŋuma* in (1a/c), it is called nominative case (NOM); when O and S have the same form, such as *ŋana* in (1b/c), it is called absolutive case (ABS). Dyirbal is said to have an NP Split Ergative system, in which the split between ergative-absolutive and nominative-accusative systems is governed by the NP type.

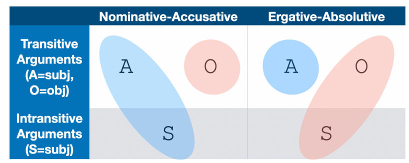

Figure 1: Schematic of ergative marking (from Papadimitriou et al., 2021), . Languages like Dyirbal have Nominative alignment (left) in 1st/2nd person but ergative in third person.

NP split ergative systems fall into an interesting typological pattern first noticed and analyzed in detail by Silverstein (1976/1986). In a Dyirbal-type split between 1st/2nd person pronouns versus all 3rd person pronouns and nouns, it is always the 1st/2nd pronouns that are on the accusative system and the 3rd person on the ergative system, never the other way around. If the split is between all pronouns and all nouns, then the pronouns are on the accusative system and nouns on the ergative system, never the other way around. More generally, splits are governed by the Nominal Hierarchy shown in Fig. 2: NP types at the left end of the hierarchy favor the accusative system while types at the right end favor the ergative system. In this paper we propose an explanation for this cross-linguistic pattern, based on an account of syntax evolution rooted in reinforcement learning.

The explanation takes off from an observation about the frequency with which the different NP types appear as subject versus object of transitive clauses (Dixon, 1994, 83ff). Every transitive clause has a subject and an object, and it is a straightforward matter to compare the number of tokens of a given NP type in subject and object functions within the transitive clauses in a corpus. For example, within transitive clauses, pronouns are more likely to function as subjects than as objects. Referring to an NP Hierarchy like Fig. 2, Dixon (1994, 85) writes that 'Those participants at the left-hand end of the hierarchy are most likely to be agents, to be in the A function, and those at the right-hand end of the hierarchy are

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

most likely to be patients, to be in the O function.'[1] Dixon continues, making a point about economy: 'It is plainly most natural and economical to 'mark' a participant when it is in an unaccustomed role.' A given NP type gets case-marked in its rarer function, so the NP case split appears to maximize the number of unmarked NPs. More generally, this property reflects a general preference towards marking the less frequent member of a pair (Haspelmath, 2021).

As shown in Fig. 2, in some language histories, accusative case starts at the left end of the hierarchy and spreads rightward; in others ergative case starts at the right end and spreads leftward (Garrett, 1990). In this paper we explain the direction of spread across the nominal hierarchy in a new way. We focus on ergative case but as noted below the same account can apply to accusative case.

We posit that the preference for NP split ergative systems (and marking of rare items more generally) is an example of communicative efficiency (Gibson et al., 2019; Haspelmath, 2021). The idea is that the marking system emerges in order to maximize effective information transmission, with minimal effort. The efficient basis of case marking has been widely explored (see Aissen, 1999, 2003; Levshina, 2021), and studied through artificial learning paradigms (Culbertson, 2012; Fedzechkina, Jaeger, & Newport, 2012; Tal, Smith, Culbertson, Grossman, & Arnon, 2022), natural language experiments (Kurumada & Jaeger, 2015), corpus analysis (Lee, 2024), and artificial neural simulation (Lian, Bisazza, & Verhoef, 2025).

Grammatical case morphemes emerge through the reanalysis of semantic case morphemes. The morpheme loses its semantic role value just when its NP serves a particular grammatical function, and becomes a marker of that function instead. Historical studies such as Garrett (1990) reveal that the linguistic condition typically preceding the emergence of NP split ergativity is one in which some arguments bearing instrumental case (INST) are subjects of transitive sentences: viz., *The key*-INST *opened the door*. Speakers reinterpret those instrumental markers as subject-of-transitive markers: viz., *The key*-ERG *opened the door*. Instruments are usually inanimate and so these new ergatives lie at the far right edge of the Nominal Hierarchy (Fig. 2), efficiently marking the rarest NP category to be subjects of transitives, namely inanimate NPs.

Thus, a key factor in the communicative efficiency of marking systems is the relative frequency of nominals in different argument positions. Silverstein (1976/1986) saw the nominal hierarchy governing splits in case-marking systems as reflecting 'semantic naturalness' (Silverstein, 1986, 174). Dixon (1994) specifically characterized the hierarchy in terms of the *likelihood* of a nominal of a given type appearing as subject versus object. Some quantitative corpus studies had by then already confirmed Dixon's view for at least some aspects of the split, in selected languages. For example,

[1] Dixon's hierarchy has a little more detail in that he separates 1st from 2nd person pronouns.

rarer as OBJ				**rarer as SUBJ of transitive**
$1^{st}/2^{nd}$ person >	3^{rd} person pron./dem. >	human nouns >	animate nouns >	inanimate nouns
OBJ case (*Accusative*) ⟶				⟵ SUBJ case (*Ergative*)

Figure 2: The Nominal Hierarchy. Case systems distinguish subject from object by marking one with a case morpheme. The morpheme appears on the *rarer* grammatical function for each NP type: pronouns are rarer as objects, so accusative pronouns tend to function as objects; lexical NPs are rarer as subjects, so ergative lexical NPs tend to function as subjects.

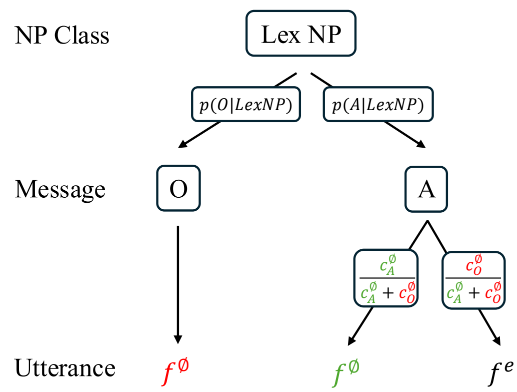

Figure 3: Schematic of the model set-up for LexNP. A speaker chooses a Noun and whether it is an A or O argument. Then, if an A argument, the speaker chooses marking in proportion to prior counts. Crucially, $p(O|LexNP)$ is empirically greater than $p(A|LexNP)$, so the high frequency of $f^{\emptyset}$ (red) means that the count $c_O^{\emptyset}$ (red) grows faster than $c_A^{\emptyset}$ (green). The same is not true for the Pronoun case (not shown).

Du Bois (1987) showed that pronouns favor the subject function while lexical NPs favor the object function in Sacapultec Maya. With the growth of corpus methods came more quantitative confirmation of the characterization in terms of likelihood (Du Bois, Ashby, & Kumpf, 2003; Haig & Schnell, 2016; Mahowald, Papadimitriou, Jurafsky, & Futrell, 2021).

Methods

What leads speakers to choose more efficient alternatives, and how exactly does that bring about language change? We addressed that question through simulations using REINFORCEMENT LEARNING (RL), using a simple learning/production algorithm, similar to those in Harley (1981) and Roth and Erev (1995). In particular, we use a method similar to the "obverter model" (Oliphant, 1996; Spike, Stadler, Kirby, & Smith, 2017), whereby a speaker decides on a form based on the distribution of meanings expressed by previous productions of that form.

Our goal is to simulate the emergence of a split ergative system, using a reinforcement learning/production setup. The basic set-up (see Fig. 3) is that a speaker who is using a LexNP (or a pronoun) chooses to use it either as an A (transitive agent) or an O (transitive object). We assume they do so in proportion to the frequency with which a LexNP (or a pronoun) appears in A or O position (which is something we estimate empirically as $P(A|LexNP)$ and $P(O|LexNP)$). Crucially, across languages LexNPs are more likely to be O's than A's, while Pronouns are more likely to be A's than O's (an important asymmetry!).

If the LexNP has been chosen by the speaker to appear in O position, we assume null marking. If in A position, the speaker has to make a choice as to whether to use a null form or a marked ergative form. Our key model assumption is that the speaker simply makes this choice in proportion to how often the null form has been used in A position before, relative to O: $\frac{c_A^{\emptyset}}{c_A^{\emptyset}+c_O^{\emptyset}}$. Think of this as "given a null form LexNP, empirically how likely is it to be A as opposed to O?". This quantity depends on how often A's are realized in null form, but also critically on how often LexNP has been seen as an O. In our model, every time the speaker makes this choice, counts are updated. *All O's are unmarked*, and the use of unmarked LexNP in the A function is discouraged by the many unmarked LexNPs in the O function. As a result, the count of unmarked LexNP used as A grows more slowly than those used as O. Over time the unmarked LexNP in A can disappear entirely, leaving the marked ergative LexNPs as the only option for A.

The nominal hierarchy and empirically obtaining initial conditions A key feature of our model is how often a particular class of nominal (e.g., 1st person pronouns, animate nouns, etc.) is used as an O as opposed to an A. For convenience, we focused on Lexical NPs vs. Pronouns (which fall on distinct halves of the nominal hierarchy). We conducted our own quantitative investigation of pronouns versus lexical NPs (all non-pronouns) by using the Universal Dependencies treebanks (Nivre et al., 2020). We counted pronominal subjects, pronominal objects, lexical NP subjects, and lexical NP objects, across 169 languages. We did this for two reasons: first to test the hypothesized asymmetry and second to get estimates of plausible values for how often the lexical NPs and pronouns of a language appear in each position.

Fig. 4 shows the results for the 100 languages with the largest sample sizes. As shown, we empirically confirmed both predicted effects. First, Lexical NPs are more likely to be in O than in A. Each orange dot represents the Lexical NP counts in one language, its position showing the number of subject-of-transitive and object lexical NPs, on a log scale. Second, pronouns are more likely to be in A than in O. Each blue dot represents pronoun counts in one language. The black diagonal line represents equal counts in A and O; dots to the lower right of that line have higher counts in O than A, while dots to the upper left have higher counts in A than O. The asymmetry is confirmed across languages.[2] We use these estimates to set $P(O|LexNP)$ and $P(O|Pronoun)$ for a representative set of languages. We do not estimate things like animate vs. inanimate frequencies but note that, if there are similar frequency asymmetries, the model would make the same predictions for those. Thus, we consider LexNP vs. pronouns a generalizable instance of this phenomenon.

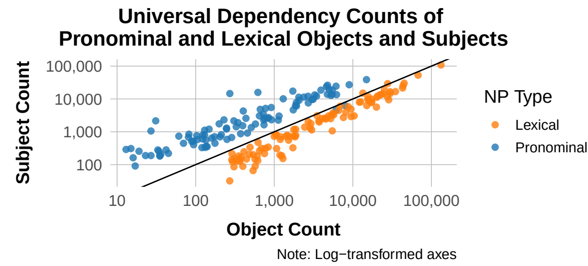

Figure 4: For each of 100 languages in Universal Dependencies, we plot the count of subjects and the count of objects, broken up by full NPs (orange) vs. pronouns (blue). Across languages, pronouns are relatively more common in subject position than they are in object position.

Reinforcement learning model In reinforcement learning, behaviors with positive outcomes are reinforced (Harley, 1981; Oliphant, 1996; Roth & Erev, 1995; Spike et al., 2017). If there are n different expressions to choose from, with previous counts $c_1 \ldots c_n$, then the probability of choosing form f_i next time is:$p(f_i|c) = \frac{c_i}{\sum_j c_j}$. After each utterance, the counts are updated. If there is an underlying preference, a prior that favors some signs over others for a particular function, then this formula has a "rich get richer" effect, in that frequent expressions get further and further reinforced. Wechsler, Shearer, and Erk (2025) use variants of the Harley-Roth-Erev formula in a model of grammar learning from the ground up that starts from nothing but a lexicon and some semantic preferences of speakers: preferences about what kinds of messages they prefer to convey, for example a preference, when talking about an action, to mention the agent.

Present study In the present study, we use this approach to explain why the pattern of ergatives and accusatives in NP split ergative languages matches the frequencies with which different kinds of NPs appear as subjects and objects.

While we expect these results to generalize across the nominal hierarchy, we here focused on lexical NPs vs. pronouns, and we modeled the emergence of ergative case with two toy languages, one for ergative lexical NPs and the other for ergative pronouns. The first toy language represents lexical NPs with the word *key* and has three sentences, shown in Table 1.

[2] Pronoun counts include affixal and independent pronouns. We did not include null pronouns, because case morphology requires an overt host. However, we collected estimates of null subject pronouns.

sentence	message	form
a. I have key.	O	$f^{\emptyset}$
b. Key opens it.	A	$f^{\emptyset}$
c. Key-ERG opens it.	A	f^{e}

Table 1: LexNP language. O is a message in which *key* is the object. A is a message in which *key* is a transitive subject. $f^{\emptyset}$ is the unmarked or zero form (*key*); f^{e} is the ergative case form (*key*-ERG).

As in Fig. 3, a speaker of the LexNP language prepares to make an utterance by first selecting a message, either message O (Table 1a; henceforth (a)) or message A (Table 1b,c; henceforth (b) and (c)), proportional to $p(O|LexNP)$. If they select O they utter (a). If they select A they utter either (b) or (c), a choice determined by the production algorithm in the next section. There are three important stages in the emergence of ergative case. At Stage I the language has only (a) and (b), with zero form ($f^{\emptyset}$) for both subjects and objects (see Table 2). At Stage II a semantic case marker on Lexical NP of transitives is reanalyzed as ergative case (ERG, f^{e}), and (c) enters the language, an alternative to (b). Then (c) becomes increasingly common until (b) has completely dropped out of the language at Stage III. In our modeling, we assume that we are already in Stage II. The production model below takes the language from Stage II to III.

The other toy language is used for modeling the emergence of ergative case on pronouns. It is the same as the LexNP language except that a pronoun replaces the word *key*. In our model, the crucial difference is that speakers of Pronoun Language select their messages proportional to $p(O|Pronoun)$, not $P(O|LexNP)$.

The production algorithm generates a different set of language histories depending on the probability distribution. Below we present the results of simulations using values from natural language data.

To do that, we use a model that keeps track of counts of the unmarked zero form $f^{\emptyset}$ and the marked ergative form f^{e}, appearing as subject in past utterances. These counts are represented by $c_A^{\emptyset}$ and c_A^{e}, respectively, where the subscript A indicates subject-of-transitive. Learners also keep track of the unmarked object forms with the count $c_O^{\emptyset}$, because objecthood competes with subjecthood as an interpretation of the unmarked form. In our learning/production model, a speaker uses these counts in the following formulae, when they must choose between producing $f^{\emptyset}$ or f^{e} for the subject (A) of a transitive clause:

(2) $$p(f^{\emptyset} \mid A) = \frac{c_A^{\emptyset}}{c_A^{\emptyset} + c_O^{\emptyset}}$$

(3) $$p(f^{e} \mid A) = 1 - p(f^{\emptyset} \mid A) = \frac{c_O^{\emptyset}}{c_A^{\emptyset} + c_O^{\emptyset}}$$

We assume that the speaker keeps track of these counts separately for two different types of NP, pronouns and lexical NPs. For simplicity, we assume that speakers observe all speech acts in the community, so that we only need to keep one set of counts for the whole speech community.

The learning/production algorithm for the LexNP Language is deployed after we already have a system in Stage II—A can be unmarked or marked, O is unmarked.[3] (The production algorithm for the Pronoun Language is the same except that in Step 1 the pronoun distribution is used instead of the LexNP distribution.) There are three steps:

Step 1 The speaker selects a message, A or O, proportional to $p(O|LexNP)$ (or, equivalently, $p(A|LexNP)$ since this is exactly $1 - p(O|LexNP)$).

Step 2 If they selected message O, then they use $f^{\emptyset}$, as this is their only option. They utter sentence (a) from Table 1, update the counts by adding 1 to $c_O^{\emptyset}$ (see just below), and then return to Step 1 and repeat. If they selected message A, they choose between forms $f^{\emptyset}$ and f^{e}, using these counts of previous utterances they have witnessed (including their own earlier utterances). The learning formulae produce a probability distribution over $f^{\emptyset}$ and f^{e}, from which speakers with message A select a form, and utter (b) or (c) from Table 1.

Step 3. Update the counts: if (b) (= $f^{\emptyset}$) was uttered add 1 to $c_A^{\emptyset}$; if (c) (= f^{e}) was uttered add 1 to c_A^{e}. Then return to Step 1 and repeat.

Forgetting In some of our simulations, we also model 'forgetting' with a *forgetting factor* whereby speakers 'forget' less recent speech input. We will see that forgetting speeds up convergence since the input is sampled closer to the target form. This was achieved simply by multiplying the counts c by a forgetting factor ν (*nu*), $0 < \nu \leq 1$, when updating at each new utterance. The update scheme in Step 3 of the production algorithm above is replaced with the following:

(4) Updating with forgetting, when the kth utterance has a subject in ergative case:

$$c_k^{e} = \nu c_{k-1}^{e} + 1$$

$$c_k^{\emptyset} = \nu c_{k-1}^{\emptyset}$$

The value of ν represents the strength of forgetting and the smaller ν is, the faster forgetting occurs. We used $\nu = 0.99$ for 'moderate forgetting' in the bottom-right simulation of Fig. 5. (When $\nu = 1$ we have no forgetting and we recover the original model.) Forgetting speeds up convergence, and it has its greatest effect on language changes that would otherwise be very slow. Forgetting the distant past and focusing on more recent input improves learning: it would be hard for a modern speaker to learn English from an input randomly sampled from language going back to Proto-Indo-European.

Numerical simulations We carried out numerical simulations using sample natural language data, to approximate

[3] The production algorithm is the researcher's model of human behavior. In Step 1 the speaker's selection of a message is characterized by message probabilities that we estimated in a corpus study and expressed with a probability distribution.

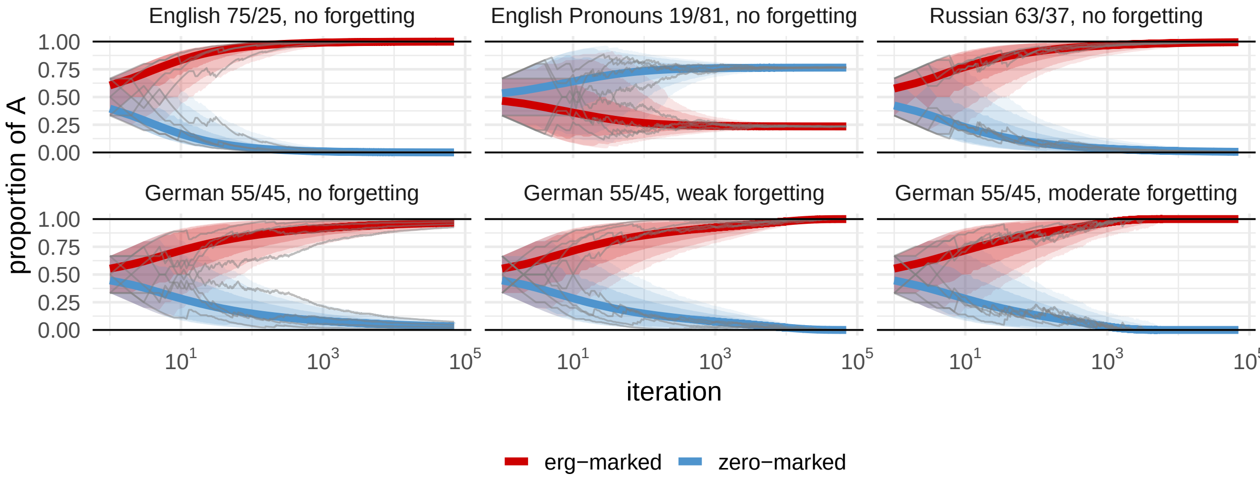

Figure 5: Simulated language histories for various settings of initial conditions based on observed frequencies in English, Russian, and German. Red represents proportion ergative case, blue unmarked. The thick lines represent the mean across simulations; shaded areas represent percentiles (.1% to 99.9%; 1% to 99%; 10% to 90%); faint lines represent randomly sampled trajectories. The top-middle graph shows the model parameter set at $p(O|Pronoun) = .19$, resulting in convergence to a steady state of about 25% ergative / 75% unmarked. The bottom row shows 3 settings of forgetting for German-like settings.

	unmarked ($f^{\emptyset}$)	ERG (f^{e})
Stage I	A $\vee$ O	
Stage II	A $\vee$ O	A
Stage III	O	A

Table 2: Ergative case becomes obligatory on subjects (A).

the speed of convergence. For each test, we computed 10,000 language histories out to 1,000,000 utterances, varying $p(O|LexNP)$ as a parameter. For $p(O|LexNP)$ we used sample values retrieved from the Universal Dependency corpus. The mean value of $p(O|LexNP)$ for all 169 languages in our sample was 0.62. We chose languages whose values of $p(O|LexNP)$ are high (English: 0.75), medium (Russian: 0.63), and low (German: 0.55). Note that the inputs are merely $P(O|LexNP)$—we derive the estimates for that quantity empirically for these languages since they vary meaningfully in this quantity, but we do not otherwise use anything specific to the languages.

We also ran simulations using the German counts with a forgetting factor. Finally, we simulated the emergence of ergative case on an NP type with a less than .5 probability of being an object. Specifically, we used the case of pronouns in the English data, where $p(O|Pronoun) = .19$.

The output of a simulation is a language history consisting of a sequence of utterances of sentences (a), (b), and (c) from Table 1.

Results

LexNP languages converge, proportional to $p(O|LexNP)$ As shown in Fig. 5, all of the simulations of lexical NPs converged to a probability of 1 for ergative case (shown in red) and zero for unmarked case (shown in blue). The speed of convergence varies by language: specifically, the higher the $p(O|LexNP)$ value, the faster the convergence. The simulation using English data ($p(O|LexNP) = .75$) is the fastest: after only 1,000 utterances, all language histories are quite close to a probability of 1. The Russian simulation ($p(O|LexNP) = .63$) is slower, and the German data simulation ($p(O|LexNP) = .55$) slower still. Convergence is slowest when lexical NPs are almost evenly divided between subject and object function, as in the German data simulation where 45% are subjects and 55% are objects.

Pronoun Languages do not converge The English Pronoun Language is one in which $p(O|Pronoun) = .19$ is low relative to $p(A|Pronoun) = .81$. The result is that, in our model, we do not see convergence to 1 in our simulations. Rather, both kinds of marking (null or marked) remain in variation on A. See the top-middle simulation in Fig. 5.

Forgetting speeds up convergence The bottom row in Fig. 5 shows the effect of forgetting in a German-like language (one with where $P(O|LexNP)$ is relatively low): forgetting speeds up convergence relative to the matched condition without forgetting.

Discussion

Our simulations show that ergative case rises in frequency and becomes obligatory on lexical NPs, in languages where lexical NPs favor objecthood over subjecthood. This condition in fact holds for lexical NPs in human languages, as we confirmed in our study of 169 languages from the Universal Dependency corpora (see Fig. 4).[4] Therefore, we have a model of the emergence of ergative case on lexical NPs in human languages.

Secondly, we found that convergence is faster if lexical NPs favor objecthood more strongly. Importantly, this holds more generally for any type T of nominal, such as the five types shown in Fig. 2. For any type T, the higher the value of $p(O|T) > 0.5$, the faster that ergative (subject) case becomes obligatory, within the model. In Fig. 2 the nominal types are arrayed from the highest $p(O|T)$ on the right to the lowest $p(O|T)$ on the left. We therefore predict that ergative case starts at the right edge on inanimate subjects and 'spreads' to the left over time. That is exactly what typological studies of human languages have shown since Silverstein's (1976/1986) original observation. Since ergative case in NP splits begins on inanimate nouns, the source for the case morpheme is generally an instrumental case marker, as observed by Garrett (1990). Instrumental semantic roles are a common role type for inanimate subjects of transitive clauses.

What is the intuition behind this account of the rise of ergative case? Why does it work? Ergative case in NP splits first arises when instrumental case is reanalyzed as ergative. In our learning model this new ergative case provides an attractive alternative to the unmarked subjects. Recall that the unmarked form is dispreferred for A because of the many unmarked forms in O (see equation (2)). Utterances of the new ergative form depress the count of unmarked subjects ($c_A^{\emptyset}$). Speakers learning from an input with a smaller proportion of unmarked forms are less likely to produce the unmarked form. They are therefore more likely to produce the only alternative, the ergative form. Thus the ergative case marker becomes more probable over time. Summarizing, the ergative case learning formula has a 'rich-get-richer' effect: ergative case use rises over time.

The strength of that rich-get-richer effect depends upon the overall likelihood that NPs of the given type (e.g. lexical NP or pronoun) are objects as opposed to subjects. A stronger rich-get-richer effect means faster convergence, which was shown to account for the Nominal Hierarchy generalization.

A third result is that ergative case on an NP type that favors subjecthood, such as pronouns, rises in probability and levels off at a probability less than one. This means the ergative case is predicted to become optional, where the probability corresponds to the frequency of use. The top-middle example in Fig. 5 shows the case where the model parameter was set at $p(O|Pronoun) = .19$ (or approximately $\frac{1}{5}$), and the resulting frequency of unmarked forms stabilized at approximately .75 (or $\frac{3}{4}$). Wechsler et al. (2025) showed through analytic means that with parametric probability p, as the number of utterances grows large, the probability of choosing the *unmarked* form converges to $(1-2p)/(1-p)$. In our example, $p = \frac{1}{5}$ of pronouns are objects. Then the predicted frequency of unmarked pronouns is $(1-2(\frac{1}{5}))/(1-\frac{1}{5}) = \frac{3}{4}$, which is what we found.

Our account of the emergence of ergative case on selected subjects can also be applied to the emergence of accusative case on selected objects. Pronoun case is the mirror image of lexical NP case: object pronouns are rarer than subject pronouns, so accusative case emerges on pronouns first (see Fig. 2). When a language has both accusative case (on NP types from the left end of the hierarchy) and ergative case (on NP types from the right end of the hierarchy), the result is a split ergative system. If accusative case appears on (object) pronouns, this is predicted to effectively block ergative from spreading to (subject) pronouns. We saw that the growth of the ergative depends on a rising count of zero form objects ($c_O^{\emptyset}$). If accusative case has replaced zero case for object pronouns, then $c_O^{\emptyset}$ cannot rise, and since $c_A^{\emptyset}$ continues to grow, it follows that ergative case will not take hold on pronominal subjects in a language with accusative object pronouns.

This account of NP splits in case marking applies only to languages in which the case inflections emerge through the reanalysis of semantic case, as in Garrett (1990). As a result it does not make absolute predictions about ergative case systems, since some systems have a different origin. For example, ergative case systems came about in some languages when the passive voice was reanalyzed as active, and the 'by-phrase' was reanalyzed as ergative case.

Conclusion

Silverstein (1986, 174) saw the nominal hierarchy as an instance of the more general observation that 'languages in general do show a relationship between surface morphological patterns and syntactic distributions on the one hand, semantic classes on the other hand.' Silverstein's classic paper, like subsequent work that it inspired, attempts to explain that relationship by modeling it directly within the grammar formalism, for example by using systems of semantic markedness features.

In this paper we have provided an explanation based on applying the version of reinforcement learning from Wechsler et al. (2025) to languages with an uneven frequency distribution of NP types across subjects and objects. The NP hierarchy emerges because the semantic class of an NP influences its frequency distribution. Our account required the minimal assumption that speakers are influenced by the observed distribution of sentence forms uttered so far. The emergent case-marking systems were furthermore shown to maximize efficiency, in the sense that they minimize morphological complexity while maintaining expressiveness.

Acknowledgments

We thank Richard Futrell for identifying the relevance of the obverter model to this work. We thank several excellent anonymous CogSci reviewers for helpful suggestions.

[4]There are different possible explanations why the subject-occurrence bias would be different for pronouns and lexical NPs. It could be discourse related: Subjects tend to be pronominal because subjects and pronouns both tend to be discourse topics. Also, pronouns tend to be animate, while lexical NPs may be animate or inanimate, and animate entities are more frequently subjects, in particular subjects of actions.

References

Aissen, J. (1999). Markedness and subject choice in optimality theory. *Natural Language & Linguistic Theory*, *17*(4), 673–711.

Aissen, J. (2003). Differential object marking: Iconicity vs. economy. *Natural language & linguistic theory*, *21*(3), 435–483.

Culbertson, J. (2012). Typological universals as reflections of biased learning: Evidence from artificial language learning. *Language and Linguistics Compass*, *6*(5), 310–329.

Dixon, R. (1994). *Ergativity*. Cambridge University Press.

Du Bois, J. W. (1987). The discourse basis of ergativity. *Language*, 805–855.

Du Bois, J. W., Ashby, W. J., & Kumpf, L. E. (2003). *Preferred argument structure*. Amsterdam: John Benjamins Publishing Company.

Fedzechkina, M., Jaeger, T. F., & Newport, E. L. (2012). Language learners restructure their input to facilitate efficient communication. *Proceedings of the National Academy of Sciences*, *109*(44), 17897–17902.

Garrett, A. (1990). The origin of NP split ergativity. *Language*, 261–296.

Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How efficiency shapes human language. *Trends in Cognitive Sciences*, *23*(5), 389-407.

Haig, G., & Schnell, S. (2016). The discourse basis of ergativity revisited. *Language*, 591–618.

Harley, C. B. (1981). Learning the evolutionarily stable strategy. *Journal of theoretical biology*, *89*(4), 611–633.

Haspelmath, M. (2021). Explaining grammatical coding asymmetries: Form–frequency correspondences and predictability. *Journal of Linguistics*, *57*(3), 605–633.

Kurumada, C., & Jaeger, T. F. (2015). Communicative efficiency in language production: Optional case-marking in japanese. *Journal of Memory and Language*, *83*, 152–178.

Lee, H. (2024). Cue reliability, communicative efficiency, and differential subject marking: Evidence from korean. *Language*.

Levshina, N. (2021). Communicative efficiency and differential case marking: A reverse-engineering approach. *Linguistics Vanguard*, *7*(s3), 20190087.

Lian, Y., Bisazza, A., & Verhoef, T. (2025). Simulating the emergence of differential case marking with communicating neural-network agents. *arXiv preprint arXiv:2502.04038*.

Mahowald, K., Papadimitriou, I., Jurafsky, D., & Futrell, R. (2021). Accessibility-based constraints on morphosyntax in corpora of 54 languages. *CUNY Sentence Processing Conference 2021*.

Nivre, J., De Marneffe, M.-C., Ginter, F., Hajič, J., Manning, C. D., Pyysalo, S., ... Zeman, D. (2020). Universal dependencies v2: An evergrowing multilingual treebank collection. *arXiv preprint arXiv:2004.10643*.

Oliphant, M. (1996). The dilemma of Saussurean communication. *BioSystems*, *37*(1-2), 31–38.

Papadimitriou, I., Chi, E. A., Futrell, R., & Mahowald, K. (2021). Deep subjecthood: Higher-order grammatical features in multilingual bert. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume* (pp. 2522–2532).

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, *8*(1), 164–212.

Silverstein, M. (1976). Hierarchy of features and ergativity. In R. M. W. Dixon (Ed.), *Grammatical categories in australian languages* (pp. 112–171). Canberra: Australian Institute of Aboriginal Studies.

Silverstein, M. (1986). Hierarchy of features and ergativity. In P. Muysken & H. van Riemsdijk (Eds.), *Features and projections* (pp. 163–232). Dordrecht: Foris.

Spike, M., Stadler, K., Kirby, S., & Smith, K. (2017). Minimal requirements for the emergence of learned signaling. *Cognitive science*, *41*(3), 623–658.

Tal, S., Smith, K., Culbertson, J., Grossman, E., & Arnon, I. (2022). The impact of information structure on the emergence of differential object marking: an experimental study. *Cognitive Science*, *46*(3), e13119.

Wechsler, S., Shearer, J. W., & Erk, K. (2025). The emergence of grammar through reinforcement learning. *arXiv preprint arXiv:2503.01635*.