

The role of gender in the acquisition of the Serbian case system

Stefan Pophristic & Kathryn Schuler*

Abstract. Serbo-Croatian is marked for seven cases and has a noun class vs. gender distinction. Given the complexity of the inflectional system, we look at Serbo-Croatian as a case study in case acquisition. We explore different correlations available in the input that children could leverage to acquire the case system in Serbo-Croatian. We ask three main questions: 1) does a noun’s gender predict the noun’s nominative singular suffix? 2) does a noun’s nominative singular suffix predict the noun’s gender? and 3) does a noun’s noun class predict the noun’s gender? Specifically, we ask whether the language input provides children with sufficient evidence to form these three productive generalizations. To test this, we apply the Tolerance Principle (Yang, 2016) to a corpus of 270 inflected Serbian nouns. Within this set of data, we find that: 1) all nominative singular suffixes productively predict a gender; 2) all genders productively predict a nominative singular suffix (with the exception of the neuter gender which predicts two suffixes); and 3) two of the three noun classes predict a single gender. We conclude that the input provides sufficient evidence for these productive correlations and we argue that children can leverage these generalizations to infer the declension patterns or gender of novel nouns. We discuss how, given these findings, children could acquire most of the inflectional system by focusing on gender as a categorization system for nouns, without needing to posit abstract categories of noun class.

Keywords. Serbian; Croatian; case marking; morphology; language acquisition

1. Introduction. Cases occur in many of the world’s languages, yet how case inflectional systems are acquired is not well understood. Slavic languages are a good subject for the study of acquisition of case, as they mark for up to seven cases and allow for word order scrambling – which makes an understanding of case inflection imperative for understanding who did what to whom. Acquiring the case system, then, is a crucial component in the acquisition process. However, the manner in which children acquire this system in Slavic languages is still an open area of research. Most research to date has focused on the diminutive advantage – the finding that caregivers use more diminutives in speech to young children, which are thought to help children acquire and use complicated case inflectional systems (see Ionin & Radeva-Bork 2017 for an overview). However, the diminutive advantage alone does not explain the acquisition process. Children who are not exposed to diminutives in their input should, in principle, still be able to acquire the whole inflectional system. Further, in many languages, including certain dialects of Serbian, diminutives occur only rarely in speech to young children. For example, Kempe and colleagues (2009) found that, for Serbian-speaking mothers, only 5% of noun tokens were introduced as diminutives in their child-directed speech, compared to 55% of nouns for Russian- and Polish-speaking mothers.

* Thanks to members of the Child Language Lab, Dr. Charles Yang, Dr. David Embick, and the University of Pennsylvania’s University Scholars program for helpful comments and suggestions and to Zorica Torbica, our monolingual native speaker. Authors: Stefan Pophristic, University of Pennsylvania (spophri@sas.upenn.edu) & Kathryn Schuler, University of Pennsylvania (kschuler@sas.upenn.edu).

In the present study, we look at Serbian (Serbo-Croatian, Bosnian, Croatian, Montenegrin, BCS) as a case study in case acquisition. All Serbo-Croatian nouns are marked for one of three genders and one of seven cases, which have a singular/plural distinction and whose suffixes are determined by one of three noun classes (see Table 1). To clarify, given varying uses of this terminology across different subfields, we adopt Kramer's (2015) definition in which gender is a nominal classification system determined by noun-external properties (which are marked in the syntax through agreement) and noun classes dictate noun-internal properties (i.e. determine a noun's specific inflection pattern). In Serbian, the three genders are: masculine (masc), feminine (fem), and neuter (neut); the three noun classes are Class I, Class II, and Class III.

Singular	Class I	Class II	Class III
Nominative	-ø/-e/-o	-a	-ø
Genitive	-a	-e	-i
Dative	-u	-i	-i
Accusative	-ø/-e/-o/-a	-u	-ø
Locative	-u	-i	-i
Instrumental	-om/-em	-om	-ju/-i
Vocative	-e/-o/-u	-o	-i
Plural			
Nominative	-i	-e	-i
Genitive	-a	-a	-i
Dative	-ima	-ama	-ima
Accusative	-e	-e	-i
Locative	-ima	-ama	-ima
Instrumental	-ima	-ama	-ima
Vocative	-i	-e	-i

Table 1. Serbo-Croatian case suffixes. The three noun classes which determine the inflectional morphology are labelled as Class I, II, and III. All 7 cases, and their respective singular and plural suffixes are given. Multiple suffixes within a cell are allomorphs (some phonologically conditioned, others not). Table adapted from Weisser (2006) and Brown & Alt (2004).

Based on findings by Marquis and Shi (2015), we know that preverbal infants already start decomposing roots and suffixes. It has also been shown that Serbo-Croatian children can use all seven singular cases from a very early age, even before the age of 2;0 (Kovačević et al. 2009). Therefore, we can assume that, by this age, children recognize that roots and suffixes can be split apart, and have at least some knowledge about which suffixes are applied in which circumstances. However, in order to fully master the case system in Serbo-Croatian, the child must learn: that these suffixes represent a case inflectional system; how many cases exist in the system; that different nouns belong to different noun classes; how many noun classes exist; and that different noun classes determine different suffixes for the same cases. All of these conclusions must be drawn from the child's input – but not all nouns will be attested in the input with all seven cases. The acquisition of this inflectional paradigm is further complicated by the

considerable presence of both homophony and syncretism across noun classes and cases (Weisser 2006). That is, a child must postulate that a suffix appearing in context A with one set of nouns may also appear in context B with a mutually exclusive set of nouns, without any presupposed notions that different sets of nouns and different contexts exist.

Despite all the complexities in the system, we propose that there are correlations between different aspects of the system in the input – noun classes, case suffixes, and gender – and that children can exploit these correlations to acquire case in Serbo-Croatian. To test this proposal, in the present study, we apply a generative learning model to a corpus of Serbian nouns to evaluate these correlations in the input. Specifically, we ask three questions:

1. Does a noun's gender predict the noun's nominative singular suffix?
2. Does a noun's nominative singular suffix predict the noun's gender?
3. Does a noun's noun class predict the noun's gender?

With these questions, we ask whether the input provides children with sufficient evidence to form productive generalizations, allowing them to infer: 1) a noun's nominative singular suffix (and thus noun class) from a noun's gender; 2) a noun's gender from its nominative singular suffix; and 3) a noun's gender from its noun class. In this way, we expose *which* generalizations are available in the system, a first step in our long-term goal of understanding *how* the child might utilize these generalizations to acquire the Serbian case system.

2. Method. To address these questions, we analyzed a corpus of Serbian nouns. Given the amount of syncretism and homophony in the Serbo-Croatian inflectional system, it is not possible to identify case and noun class with certainty from a noun's inflection alone. Therefore, we selected a corpus tagged for case: a set of 270 nouns taken from the *serbian* dataset, part of the *ndl* package in R (Arppe et al. 2018). This dataset consists of all nouns from the Kostić (1999) frequency dictionary that appear at least once with every possible case inflection. To obtain the gender for each noun in the dataset, the first author, a heritage speaker of Serbo-Croatian, hand-coded each noun for gender. Then, we cross-checked these codings with a mono-lingual native speaker of Serbo-Croatian.

In our analysis, we referred to the /-ø/ suffix of the nominative singular in noun classes I and II (Table 1) by the consonant that the noun's root ends in (/C/). The null suffix is almost exclusively applied to roots that end in consonants. We did this because we could extract words ending in /C/ in the nominative singular case from our corpus, but could not do so with words ending in a suffix which is not overtly realized. This is in line with the child's input, as the child is exposed to nominative singular nouns that end in /C/, /-a/, /-o/ and /-e/. Only after this exposure, and learning the inflectional system, can the child posit the existence of a null morpheme on the words that end in /C/. Likewise, within this analysis we only focus on the singular forms of the cases, as all 7 cases in the singular are attested in child utterances before the age of 2;0, but the plural forms are not all attested even by the age of 2;9 (Kovačević et al. 2009).

In order to establish which correlations are productive for the child learner, we also need to define productivity. Here, by *productive* we mean that the input provides the child with sufficient evidence to form a productive generalization. Specifically, we will ask whether Serbian-learning children have sufficient evidence to generalize these correlations (e.g. that a noun with X gender will belong to Y noun class). We will use the Tolerance Principle (Yang 2016) to establish a threshold for productivity, as it gives explicit yet conservative predictions about which morphological rules could be learned productively by the language learner. This learning model

has been shown to correctly predict generalizations in corpus data from a wide range of different languages (e.g. Fernández-Dobao and Herschensohn 2019; Merkuur et al. 2020; Garcia 2019; Björnsdóttir, 2021) and childhood behavior in certain experimental settings (Schuler et al. 2016). We expect our findings to be compatible with other generative learning models.

The Tolerance Principle quantifies the number of exceptions a productive rule can tolerate before its formation becomes computationally inefficient. That is, it assesses when forming a rule is no longer more computationally efficient than storing each lexical item and its inflections individually. As per Yang (2016), it is defined as follows: “If R is a productive rule applicable to N candidates, then the following relation holds between N and e , the number of exceptions that could but do not follow R :

$$e \leq \theta_N \text{ where } \theta_N := \frac{N}{\ln(N)}$$

The Tolerance Principle only needs two values from the input to determine whether a rule is productive: the number of items to which the rule can potentially apply, N , and the number of exceptions to the rule, e . For example, if we wanted to posit a rule that nouns of gender X will belong to noun class Y, we would have to find the number of nouns that this rule could apply to (i.e. number of nouns with gender X), and how many of those nouns are exceptions to the rule (i.e. do not belong to noun class Y). If we have 10 nouns of gender X in our language ($N = 10$), our rule can tolerate $\theta_{10} = 4.3$ exceptions. If 3 of those nouns do not belong to noun class Y ($e = 3$), our rule will be productive ($e < \theta_n$ or $3 < 4.3$). On the other hand, if we have 6 nouns that do not belong to noun class Y ($e = 6$), then our rule will not be productive ($e > \theta_n$; $6 > 4.3$).

In order to apply the Tolerance Principle, we reformulate our 3 questions above into the following productive rules:

- Rule 1. Nouns of gender X have nominative singular suffix Y.
- Rule 2. Nouns with nominative singular suffix Y have gender X.
- Rule 3. Nouns of noun class Z have gender X

We then applied the Tolerance Principle to these correlations. Using our tagged corpus, we found N and e for each Rule 1-3. We then applied the Tolerance Principle to those values to determine whether the corpus contained sufficient evidence to form that generalization. If there is sufficient evidence, we can conclude that these correlations can be learned as productive rules by a child acquiring the language. The specific parameters (i.e. how we defined N and e) for each question are given below.

1. Does a noun’s gender productively predict the noun’s nominative singular suffix?

Generalization: *Nouns of gender X have nominative singular suffix Y*. Since we want to determine whether a specific gender predicts a nominative singular ending, we grouped the nouns based on their gender: N = the number of nouns belonging to one of the three genders. We then calculated how many of those N nouns within a single gender have a given nominative singular suffix (i.e. follow the rule; $N-e$) and how many do not (i.e. are exceptions to the rule; e).

2. Does a noun’s nominative singular suffix productively predict the noun’s gender?

Generalization: *Nouns with nominative singular suffix Y have gender X*. We grouped nouns based on their nominative singular endings: N = number of nouns that share a single nominative

singular case suffix. We then calculated how many of those N nouns have a given gender (i.e. follow the rule; $N-e$) and how many do not (i.e. are exceptions to the rule; e).

3. Does a noun's noun class productively predict the noun's gender?

Generalization: *Nouns of noun class Z have gender X* . While nominative singular suffixes are determined by noun class, both noun classes I and III have the $/-C/$ ending as a nominative singular suffix. In order to tease apart correlations between noun classes with the same nominative singular suffix, we tested the correlation between noun class and gender, using the genitive singular ending as a proxy for noun class instead of the nominative singular. Just like the nominative singular suffixes, genitive singular suffixes are determined by noun class. However, unlike their nominative counterparts, the genitive singular case markers have nosyncretism/homophony across noun classes. While the genitive case resolves the issues mentioned above, we still wanted to focus questions (1) and (2) on the nominative singular case, for the following reasons: it is the most frequent case in the child's input (Kovačević et al. 2009); it is the base case in the language; it is the first case the child utters (Kovačević et al. 2009); and it gives us interesting predictions about what children can use to aid their acquisition of this system.

For this third question, we grouped nouns based on their genitive singular ending; N = number of nouns which have the same genitive singular ending. We then calculated how many of those N nouns have a given gender (i.e. follow the rule; $N-e$) and how many do not (i.e. are exceptions to the rule; e).

3. Results. Results from applying the Tolerance Principle to our first generalization, *nouns with gender X have nominative singular suffix Y* , are shown in Table 2. By definition nominative singular ending is determined by noun class – therefore this question essentially asks whether a child can assume that a noun belongs to a specific noun class based on the noun's gender.

gender	TP		nominative singular				productive
	N	θ_n	-C	-a	-o	-e	
feminine	120	25	9	110	1	0	-a
masculine	112	23	107	0	5	0	-C
neuter	38	10	0	0	22	16	-o/-e

Table 2. Tolerance Principle calculations for question 1. For each gender, we indicate the number of nouns of that gender in the corpus (N) and the Tolerance Principle threshold for productivity (θ_n). We then indicate how many of these N nouns take each nominative singular suffix and which of these emerge as the productive rule according to the Tolerance Principle.

To illustrate our approach, consider the feminine gender. In our corpus of 270 nouns, 120 of them were feminine (N). According to the Tolerance Principle, in order for a productive rule to be formed, $120/\ln(120) = 25$ of the feminine nouns in the corpus can be exceptions to this rule (θ_n). As shown in Table 2, 110 of nouns of gender X (feminine) take Y noun class (-a). The remaining 10 are exceptions (9 are -C and 1 is -o). Therefore, we can conclude that the feminine gender productively predicts the $/-a/$ nominative singular suffix. The $/-a/$ suffix belongs to noun class II (Table 1), therefore the feminine gender productively predicts that the noun will belong to noun class II. The masculine gender productively predicts the $/-C/$ nominative ending. However, this ending is found

both in noun class I and III (Table 1). Therefore, it remains unclear which noun class the masculine gender predicts. We will return to this issue in rule 3. The neuter gender was the only gender which did not predict any single suffix. However, a rule that predicts either /-o/ or /-e/ as a suffix for a neuter noun had no exceptions. Both of these endings belong to noun class I as nominative singular endings (Table 1). Therefore, the neuter gender productively predicts noun class I belonging.

Both the feminine and neuter genders predict noun class belonging productively. Given these rules, we can conclude that if a child knows a noun's (feminine or neuter) gender, the child can assume the noun's declension pattern. Although there are not many cases where a noun's gender is heard, but the noun's declensional suffix is not, these rules can still serve as an extra cue in determining noun class belonging for novel nouns.

For the first rule, we tested whether knowing the gender of a noun will indicate the noun's nominative singular ending. Here we are testing the converse—could a child infer the gender of a noun given its nominative singular suffix? Recall that, to apply the Tolerance Principle, we formalized this rule in Rule 2, *nouns with nominative singular suffix Y have gender X*.

nominative singular	TP		gender			productive
	N	θ_n	feminine	masculine	neuter	
-a	110	23	110	0	0	feminine
-C	116	24	9	107	0	masculine
-o	28	8	1	5	22	neuter
-e	16	5	0	0	16	neuter

Table 3. Tolerance Principle calculations for question 2. Calculations for all nominative singular ending vs. gender combinations are shown. Productive correlations are given in the rightmost column.

Within this corpus, the /-a/ suffix always predicts the feminine gender. Since /-a/ is only specified by noun class II in the nominative singular (Table 1), we can conclude that noun class II always predicts the feminine gender. The /-C/ nominative ending productively predicts the masculine gender. Again, because the /-C/ ending is found in both noun class I and III we cannot conclude from this question alone which noun class productively predicts the masculine gender. Both the /-o/ and the /-e/ nominative singular suffixes productively predict the neuter gender. Since both belong to noun class I, noun class I could productively predict the neuter gender. However, since noun class I also has the /-C/ ending as a nominative singular ending (table 1), it remains unclear which suffix (and through productive correlation, gender) noun class I predicts. These findings suggest that if a child hears a noun in the nominative singular form with 3 of the 4 possible nominative singular suffixes (/a/, /o/, and /e/), the child will be able to infer the noun's gender. This is especially helpful for the acquisition of novel nouns, as not all syntactic contexts mark for gender, but all contexts require case declination of nouns. As mentioned previously, the nominative singular is the most frequent case in the child's input. Therefore, given these findings, we predict that a child would be able to learn a novel noun from a syntactic

context that does not mark for gender (given the noun was marked one of three of the nominative singular suffixes mentioned above) and use it in syntactic contexts that do mark for gender.

The findings in questions (1) and (2) still leave us with several questions regarding the relationships between noun classes, their nominative singular suffixes, and gender. While the masculine gender and /-C/ ending are productively correlated, the /-C/ ending can belong to either noun class I or III. Therefore, we still can't predict noun class belonging solely based on the noun's masculine gender. Furthermore, we know that the masculine gender predicts the /-C/ nominative singular ending and that neuter predicts the /-o/ and /-e/ suffixes. However, all three of these nominative singular endings belong to the same noun class (noun class I). Therefore, while every nominative singular ending may predict a gender, noun class I could productively predict either masculine, neuter, or neither. Our question (3) uses the genitive singular as a proxy for noun class, rather than the nominative singular, to resolve these issues. Our calculations for every genitive singular suffix/noun class and gender are given below. Productive correlations are given in the rightmost column.

Noun Class	TP		gender			productive
	N	θ_n	feminine	masculine	neuter	
I	150	29	0	112	38	not feminine
II	112	23	112	0	0	feminine
III	10	4	10	0	0	feminine

Table 4. Tolerance Principle calculations for question 3. Calculations for all noun class/genitive singular suffix vs. gender combinations are shown.

We found that noun class I is not productive for any single gender, but productively indicates that a noun is not feminine, and that noun classes II and III are both productive for the feminine gender. This data resolves the two gaps from the findings in questions 1 and 2. As shown above, noun classes I and III both have the /-C/ nominative singular ending, which impedes us from concluding whether the masculine gender predicts any noun class. As we can see with our findings to question 3, noun class I is productive for any gender other than feminine without any exceptions and noun class III is productive for the feminine gender without any exceptions. Therefore, in the nominative singular, feminine nouns ending in /-C/ and masculine nouns ending in /-C/ are mutually exclusive in terms of noun class belonging. Since the masculine gender is a productive indicator of the /-C/ ending and masculine nouns ending in /-C/ are found in noun class I (and never in III), we can conclude that the masculine gender productively predicts noun class I belonging. On the other hand, as we see with noun class I, not every noun class predicts a single gender. While noun class I has three possible nominative singular endings, each productively correlated with either the masculine or neuter gender, none of the two genders wins out as the productive gender of the noun class.

These correlations leave us with a specific prediction that if a child hears a noun with the /-C/ ending, the child can assume that the noun is 1) masculine, and 2) belongs to noun class I rather than III. Furthermore, if the child hears a novel noun declined for a non-nominative case and the child does not know the noun's gender, the child can infer the noun's gender if the noun is declined for noun classes II and III (barring exceptions of homophony/syncretism for a case

across noun classes). The child cannot make specific inferences about the nouns gender if the noun is declined in noun class I; but they can assume that the novel noun is masculine or neuter (i.e. not feminine).

4. Discussion. Taken together, our results suggest that if a child hears a noun in the nominative form without knowing its gender, the child can use a productive generalization to determine its gender. Similarly, if a child hears a noun's gender, the child can use a productive generalization to generate the noun's declension pattern and (with the exception of the neuter) generate the noun's nominative singular ending. These productive correlations allow the child to assume declension patterns or gender for a novel noun without having heard one or the other in the input. Thus, we predict that the child would be able to take a novel noun which was heard in only one specific syntactic context and use it in novel syntactic contexts which may require overt gender marking or different case declensions.

Of particular interest is the finding that all three genders predict noun class belonging. As discussed in the introduction, children must draw several conclusions about the structure of the case inflection system despite messy input. The child could approach the acquisition process with an assumption of a distinction between noun class and gender. However, this is unlikely to emerge without positive evidence in the input, and accumulating enough positive evidence would take a prohibitively long amount of time, given how complicated the system is and how much homophony exists. However, positing a gender distinction in their grammar is simpler. Since gender is motivated through noun external evidence (through overt marking in the pronouns, certain verb tenses, determiners, etc), the child does not need any case-inflection-specific input to posit this distinction. Given the salience of the gender system, and our finding that gender predicts noun class productively, we can posit the following theoretical acquisition model. Children first group nouns into mutually exclusive sets based solely on gender. Then, using the productive generalizations supported by the input, children could learn the nouns' inflections within each of those sets. Following this method, children would almost exclusively be exposed to sets of nouns within a single noun class, and thus sets of nouns that all take the same case inflections. Learning these inflections by sorting nouns into these externally-motivated and systematic categories would cut down on any confusion caused by syncretism/homophony across noun classes, making the system easier to learn. After the child learns that the system has seven cases, what the cases are, and which suffixes belong to each gender, the child can more easily focus on irregularities of the system which would lead them to posit the existence of noun classes. Such a bootstrapping method could potentially be utilized by children acquiring languages that have both noun classes and gender (e.g. Harris 1991 for Spanish).

However, we cannot draw any such conclusions from this data alone. Our findings show that these correlations exist in the data and could theoretically be employed by the child. We hope to continue this line of research through assessing these specific predictions both in corpus data of child utterances and through elicited production in wug tests (Berko 1958). Further research on these correlations, and the potential bootstrapping method discussed above, inherently explore questions about which distinctions a child assumes when approaching the acquisition problem. Does the child assume a distinction between noun-internal (i.e. noun class) and noun-external (i.e. gender) properties before exposure to the data? If so, do we have to attribute this distinction to an innate feature of the grammar? Or does the child favor one in order to learn the system and acquire evidence to support the existence of the other. We also hope that further research of this bootstrapping method could help inform discussions about the directionality of the correlation

between gender and noun class in the adult morphology (e.g. Aronoff 1994).

5. Conclusions. We applied a generative learning model to assess the productivity of different correlations between nominative singular endings, noun classes, and gender in Serbo-Croatian. In particular, we tested: 1) whether a noun's gender predicts a noun's nominative singular suffix; 2) whether a noun's nominative singular suffix predicts the noun's gender; and 3) whether a noun's noun class predicts the noun's gender. Based off of these results, we concluded that: a child can assume the gender of a noun solely based off of its nominative form; a child can assume a noun's gender based off of a non-nominative case declension for 2 of 3 noun classes in the system; and that a child can assume a noun's declension pattern and (with the exception of the neuter gender) assume the noun's nominative singular ending solely based off of the noun's gender. Possible applications of these conclusions for the acquisition of Serbo-Croatian, and gender, noun class, and case systems in general were discussed. We hope that through continuing this research of the acquisition of the Serbo-Croatian case system, we can not only help shed light on the acquisition of this language but also the acquisition of case and the relationship between noun class and gender more generally.

References

- Aronoff, Mark. 1994. *Morphology by itself*. Cambridge, MA: MIT Press.
- Berko, Jean. 1958. The child's learning of English morphology. *Word* 14(2-3). 150–177.
- Arppe, Antti, Peter Hendrix, Petar Milin, R. Harald Baayen, Tino Sering & Cyrus Shaoul. 2018. ndl: Naive Discriminative Learning. R package version 0.2.18.
<https://CRAN.R-project.org/package=ndl>.
- Björnsdóttir, Sigríður Mjöll. 2021. Productivity and the acquisition of gender. *Journal of Child Language* First view. 1–26. <https://doi.org/10.1017/S0305000920000732>.
- Brown, Wayles & Theresa Alt. 2004. *A handbook of Bosnian, Serbian, and Croatian*. Durham, NC: Slavic and East European Language Research Center (SEELRC), Duke University. <http://www.seelrc.org:8080/grammar/mainframe.jsp?nLanguageID=1>.
- Fernández-Dobao, Ana & Julia Herschensohn. 2020. Present tense verb morphology of Spanish HL and L2 children in dual immersion: Feature Reassembly revisited. *Linguistic Approaches to Bilingualism* 10(6). 775–804. <https://doi.org/10.1075/lab.18026.fer>.
- Garcia, Guilherme D. 2020. Language transfer and positional bias in English stress. *Second Language Research* 36(4). 445–474. <https://doi.org/10.1177/0267658319882457>.
- Harris, James W. 1991. The exponence of gender in Spanish. *Linguistic Inquiry* 22(1). 27–62.
- Ionin, Tania & Teodora Radeva-Bork. 2017. The state of the art of first language acquisition research on Slavic languages. *Journal of Slavic Linguistics* 25(2). 337–366.
<https://doi.org/10.2307/26535069>.
- Kempe, Vera, Nada Ševa, Patricia J. Brooks, Natalija Mironova, Angelina Pershukova, & Olga Fedorova. 2009. Elicited production of case-marking in Russian and Serbian children: Are diminutive nouns easier to inflect? *First Language* 29(2). 147–165.
<https://doi.org/10.1177/0142723708092441>.
- Kostić, Djordje. 1991. *Frekvencijski rečnik savremenog srpskog jezika*. Beograd: Institut za eksperimentalnu fonetiku i patologiju govora.
- Kovačević, Melita, Marijan Palmović & Gordana Hržica. 2009. The acquisition of case, number and gender in Croatian. In Ursula Stephany & Maria D. Voeikova (eds.), *Development of*

- nominal inflection in first language acquisition: A cross-linguistic perspective*, 153–177. Berlin: Mouton de Gruyter.
- Kramer, Ruth. 2015. *The morphosyntax of gender*. Oxford: Oxford University Press.
- Marquis, Alexandra, & Rushen Shi. 2015. The beginning of morphological learning: Evidence from verb morpheme processing in preverbal infants. In Roberto G. de Almeida & Christina Manouilidou (eds.), *Cognitive science perspectives on verb representation and processing*, 279–295. New York: Springer. https://doi.org/10.1007/978-3-319-10112-5_13.
- Merkuur, Anne, Jan Don, Eric Hoekstra & Arjen P. Versloot. 2019. Competition in Frisian past participles. In Franz Rainer, Francesco Gardani, Wolfgang U. Dressler & Hans Christian Luschützky (eds.), *Competition in inflection and word-formation*, 195–222. New York: Springer.
- Schuler, Kathryn D., Charles Yang & Elissa L. Newport. 2016. Testing the Tolerance Principle: Children form productive rules when it is more computationally efficient to do so. *CogSci* 38.
- Weisser, Philipp. 2006. A distributed morphology analysis of Croatian noun inflection. In Gereon Müller & Jochen Tromme (eds.), *Subanalysis of argument encoding in Distributed Morphology*, 131–142. Leipzig: Universität Leipzig.
- Yang, Charles. 2016. *The price of linguistic productivity: How children learn to break the rules of language*. Cambridge, MA: MIT press.