RaY

Research at the University of York St John

Acquisition and Representation of Grammatical Categories: Grammatical Gender in a Connectionist Network

Jelena Mirković (j.mirkovic@psych.york.ac.uk)
Department of Psychology, University of York
York, YO10 5DD, UK

Mark S. Seidenberg (seidenberg@wisc.edu)
Department of Psychology, 1202 W. Johnson Street
Madison, WI 53706, USA

Maryellen C. MacDonald (mcmacdonald@wisc.edu)
Department of Psychology, 1202 W. Johnson Street
Madison, WI 53706, USA

Abstract

In traditional models of language production grammatical categories are represented as abstract features independent of semantics and phonology. An alternative view is proposed where syntactic categories emerge as a higher-order regularity from semantic and phonological properties of words. The proposal was tested using grammatical gender in Serbian, a south Slavic language with rich morphology. Semantic and phonological correlates of gender are described using a corpus of 1221 Serbian nouns. A PDP network was trained to produce the same words based on distributed semantic representation as input and distributed phonological representation as output, and with no explicit representation of grammatical gender. Upon successful learning of the training corpus, generalization was explored using test corpora designed to capture semantic and phonological properties of different genders. The findings suggest that grammatical gender, as other syntactic categories, may be viewed as emerging through coherent co-variation of semantic and phonological properties of words during learning.

Keywords: grammatical categories; grammatical gender; modeling; connectionism

Introduction

In many language production and comprehension models syntactic categories are represented as abstract symbolic features independent of semantic and phonological properties of words (e.g. Caramazza, 1997; Levelt, Roelofs, & Meyer, 1999). However, there is evidence that grammatical categories are correlated with both phonological and semantic properties of words. For example, verbs tend to denote actions whereas nouns tend to denote objects (e.g. Pinker, 1984). In English, verbs have fewer syllables than nouns (e.g. Kelly, 1992). Recently there has been increasing evidence that speakers encode this kind of information both in acquisition and processing (e.g. Cassidy & Kelly, 1991; Schwichtenberg & Schiller, 2004).

Here we aim to explore acquisition and representation of syntactic categories, specifically grammatical gender, using the PDP theoretical framework. This approch has often been used in exploring acquisition and representation of aspects of language that are considered abstract and typically treated as symbolic in traditional models of language processing (e.g. rules in reading acquistion, rules in morphological processing; see, e.g., Gonnerman, Seidenberg, & Andersen, 2007; Harm & Seidenberg, 2004, as recent examples). This research has shown that phenomena that at first sight seem difficult to explain without a recourse to symbolic representations can actually be explained within systems that encode probabilistic phonological and semantic properties of words. To a great extent the claims of this approach have relied on properties of processing in PDP networks. One of the properties of interest here is their ability to encode coherent co-variation of input-output mappings. In semantic cognition, for instance, Rogers and McClelland (2004) have shown that a simple feedforward network develops internal representations that resemble the hierarchical structure of semantic categories based on consistently more complex covariation of the input-output mappings. We hypothesize that grammatical categories can be viewed in a similar way to the extent they are correlated with semantic and phonological properties of words.

A similar approach has been used to explore acquisition of categories like nouns and verbs (e.g. Monaghan, Chater, & Christiansen, 2005; Shi, Morgan, & Allopenna, 1998), but focusing on phonological and/or distributional cues. Here we test the proposed hypothesis using grammatical gender, a syntactic category that is typically described as *semantically* arbitrary (e.g. Bates, Devescovi, Hernandez, & Pizzamiglio, 1996; Caramazza, 1997). However, in an extensive linguistic study Corbett (1991) argues that all gender systems have a semantic core based on e.g. natural gender or animacy. Furthermore, grammatical gender is correlated with more fine-grained semantic cues: e.g. in German nouns referring to alcoholic drinks tend to be masculine, whereas nouns referring to reptiles tend to be feminine (Zubin & Köpcke, 1986). This allows us to also explore semantic regularities in grammatical categories, which have typically not been explored within this approach.

The goal of the present study is two-fold. First, in a corpus analysis we explored semantic and morphophonological properties of nouns of different genders in Serbian, a south

Slavic language with a complex morphological system. This analysis aimed to establish what is available in the environment of the native speaker, i.e. to describe the structure of the input. To foreshadow the results, as in other languages, Serbian grammatical gender is correlated with both semantic and mophophonological lexical properties.

Second, we explored whether a system of this kind can be learned based solely on phonological and semantic information, and no explicit information about grammatical gender. For this purpose we used a connectionist network whose task was to produce words based on a distributed semantic representation as input and a distributed phonological representation as output. Specifically, we hypothesized that grammatical gender can be viewed as coherent co-variaton of semantic and phonological properties that emerges through learning. For example, as words with similar semantic properties (e.g. beverages containing alcohol) are being learned, their similarity in the semantic space is accompanied by their similarity in the phonological space (e.g. ending in /e/). We suggest that this co-variation in the semantics<->phonology mapping gives rise to grammatical gender, and perhaps *is* what is more conveniently termed grammatical gender. The performance of the model was analyzed to address this hypothesis. The results suggest that the model that did not have an explicit representation of grammatical gender successfully learned the system that encodes it. Furthermore, the network's performance was influenced by the co-variation in the semantics<->phonology mapping, such that the more consistently co-varying items were easier for the model in generalization tests.

Corpus Analyses

Corpus of Serbian Nouns

Serbian nouns are coded for grammatical gender (masculine, feminine, neuter), number and case within a single suffix (see Table 1 for examples). As in other Indo-European languages, Serbian gender is based on natural gender, such that in animate nouns there is a correspondence between natural and grammatical gender. Nouns agree with modifiers in gender, number and case. If they are a subject of the sentence they agree with the verb in number, as well as gender in some forms. In many of these properties Serbian is similar to other gendered languages. For example, like in German there are three genders, and there is agreement between the noun and the modifier like in Spanish and Italian. Thus we expect the findings about this language to generalize to other languages.

The corpus used in this study consisted of the Serbian translations of a subset of nouns used to develop semantic feature norms for English (McRae, Cree, Seidenberg, & McNorgan, 2005). McRae and colleagues asked participants that for a given set of words they list features of the things the words referred to. Since all items used in these norms were concrete nouns, it is plausible to assume that similar features would be produced by speakers of different languages, assuming that the concepts exist in both languages. The items where a clear

Table 1: Inflectional paradigm of a masculine (/medved/ – bear), a feminine (/krava/ – cow) and a neuter noun (/selo/ – village)

Singular forms			
case	masculine	feminine	neuter
nominative	medved	krav**a**	sel**o**
genitive	medved**a**	krav**e**	sel**a**
dative	medved**u**	krav**i**	sel**u**
accusative	medved**a**	krav**u**	sel**o**
instrumental	medved**om**	krav**om**	sel**om**
locative	medved**u**	krav**i**	sel**u**
vocative	medved**e**	krav**o**	sel**o**

Plural forms			
case	masculine	feminine	neuter
nominative	medved**i**	krav**e**	sel**a**
genitive	medved**a**	krav**a**	sel**a**
dative	medved**ima**	krav**ama**	sel**ima**
accusative	medved**e**	krav**e**	sel**a**
instrumental	medved**ima**	krav**ama**	sel**ima**
locative	medved**ima**	krav**ama**	sel**ima**
vocative	medved**i**	krav**e**	sel**a**

Note: The examples are in the International Phonetic Alphabet. The suffixes are underlined.

direct translation could not be provided were excluded (e.g *cherry* in English includes sweet and sour cherry, whereas in Serbian two different words are used).

We first performed a hierarchical cluster analysis on the features associated with each word in order to identify semantic categories in the corpus, as semantic categories were used as the basic level of analysis in studies on semantic regularities in grammatical gender (e.g. Corbett, 1991). This procedure yielded 99 masculine and 85 feminine nouns that were clustered in 15 semantic categories (e.g. animals, furniture, fruit, musical instruments). This ratio of masculine and feminine nouns is representative of the ratio obtained in larger corpora (Mirković, Seidenberg, & Joanisse, 2008). The original set of English concepts did not provide enough nouns whose translations were neuter in Serbian, so all the analyses included masculine and feminine nouns only. Since masculine and feminine nouns comprise more than 85% of nouns in Serbian a more detailed analysis of this set can be considered representative in terms of establishing the kinds of regularities that exist in the language.

Semantic Regularities

In order to establish the extent to which different semantic categories are associated with different genders, we calculated the proportion of masculine and feminine nouns in each semantic category (Figure 1).

There are only 3 categories which are highly predictive of grammatical gender, i.e. where all nouns are of one gender: vehicles (masculine), buildings (feminine), and kitchen containers (feminine). Examples of categories predictive of gender but less strongly are birds (70% masculine), fruit (69.23% feminine), and fish (67% feminine). There are also categories not predictive of gender, for example furniture (50% femi-
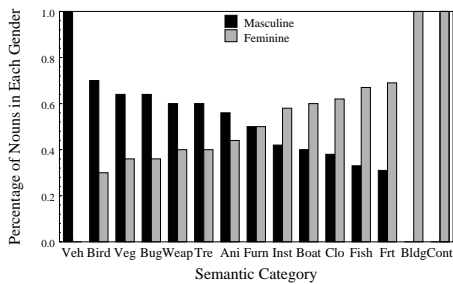
Figure 1: Semantic category membership as a cue to grammatical gender.
veh = vehicles; veg = vegetables; bug = bugs; weap = weapons; tre = trees; ani = animals; furn = furniture; inst = musical instruments; boat = boats; clo = clothing; frt = fruit; bldg = buildings; cont = kitchen containers

nine, 50% masculine) and animals (55.9% masculine, 44.1% feminine). These findings indicate that in Serbian, as in other languages, there are semantic regularities that are correlated with grammatical gender, which are not limited just to nouns referring to animate entities with natural gender.

Morphophonological Regularities

This analysis was performed on all inflectional forms of the 184 nouns, as reported in the Frequency dictionary of Serbian (Kostić, 1999), a total of 1231 words. Two analyses were performed: one focuses on the distribution of word endings in the two genders, and the other on word internal phonological properties.

Word endings The percentage of inflectional forms where word endings differ for the two genders is 70.99. Even though there is a large overlap in specific phonemes used in the two genders (see Table 1) they are associated with different inflectional forms in the two genders. As inflectional forms code for case, this implies that in the same sentential context masculine and feminine nouns are associated with *different* endings in the majority of cases. This shows that word endings are a strong cue to grammatical gender, similar to findings in other languages. However, this cue is probabilistic as in almost 30% of forms the ending is gender-ambiguous.

Word-internal phonological properties Word internal phonological properties such as syllabic structure have been shown to be correlated with grammatical class (e.g. Kelly, 1992). We focused on the distribution of phonological properties across different inflectional forms of the same word (analogous to *knife– knives* vs. *book–books* in English), and the extent to which the two genders differ in this regard.

We identified 4 groups of nouns in masculine gender. About a half of masculine nouns (54.55%) have a simple pattern where inflectional forms are derived by changing the suffix and leaving the stem unchanged (see Table 1). A quarter of masculine nouns (25.25%) gets an infix between the stem and the inflectional suffix in all plural forms,

whereas singular forms involve only a suffix change (e.g. /lav/(lion)-nom.sg., /lava/-acc.sg., /lav**ovi**/-nom.pl., /lav**ove**/-acc.pl.). In 16.16% of masculine nouns the syllabic structure changes across inflectional forms because of an epenthetic vowel, which is sometimes accompanied by other phonological alternations e.g. devoicing (e.g. /vra**bats**/(sparrow)-nom.sg., /vra**ptsa**/-gen.sg.). A small group of masculine nouns (4.04%) has palatalization of the final consonant of the stem (/jastog/(lobster)-nom.sg., /jasto**zi**/-nom.pl.).

An overwhelming majority (88.24%) of feminine nouns shows no stem changes across different inflectional forms (see Table 1). Slightly less than 10% of feminine nouns involves an epenthetic vowel which changes the syllabic structure of the word (e.g. /ovtsa/(sheep)-nom.sg., /ovatsa/-gen.pl.), and 2.35% of feminine nouns have palatalization in some inflectional forms (e.g. /jabuka/(apple)-nom.sg., /jabu**tsi**/-loc.sg.).

In summary, unlike masculine nouns, in the majority of feminine nouns there is a high degree of phonological overlap between the word's inflectional forms, i.e. only the inflectional suffix changes and the rest of the word remains intact. By contrast, in about 40% of the masculine nouns there is greater phonological dissimilarity among the word's inflectional forms: there are differences not only in the inflectional ending, but also in length (/lav/(lion)-nom.sg., /lavovi/-nom.pl.), and syllable and/or phonemic structure (e.g. /vrabats/(sparrow)-nom.sg., /vraptsa/-gen.sg.). This indicates that inflectional forms of masculine nouns have more word-internal variation across inflectional forms relative to feminine nouns.

Discussion

The findings from the corpus analyses demonstrate that Serbian nouns of different genders show various regularities in their semantic and phonological properties. Are these regularities enough to give rise to grammatical gender during learning or is an explicit feature necessary? This is one of the hypothesis tested in the computational model presented below. Furthermore, nouns of different genders also show regularities in the semantics<->phonology mapping. For example, words denoting vehicles tend to be nouns with morphophonological properties associated with masculine gender, whereas words denoting buildings and kitchen containers tend to be associated with morphophonological properties of feminine gender. When a child is learning the language this means that coherently co-varying semantic properties that give rise to semantic categories (Rogers & McClelland, 2004) are accompanied by coherently co-varying phonological properties, e.g. a particular set of inflectional endings. In addition, if there is agreement among different words in the sentence the coherent co-variation of properties extends beyond single words, e.g. to modifiers and verbs. It is possible that this coherent co-variation of different lexical properties gives rise to grammatical gender. One of the implications of this view is that coherent co-variation in the semantics<->phonology mapping (S<->P consistency) should produce fa-

Figure 2: Architecture of the model. Ellipses represent groups of units organized into layers. Arrows indicate directional weighted connections used to pass information among layers. The semantic layer is connected to the phonological output through a recurrent hidden layer with 250 nodes.

cilitative effects in processing. For instance, the items denoting buildings (semantically consistent since they belong to the same category, morphophonologically consistent since all of them have properties associated with feminine nouns) should be easier to learn in a PDP network relative to the items denoting furniture (where the similarity in the morphophonological space is significantly smaller since half of the nouns are associated with the properties of feminine nouns, and half with masculine). This is the second hypothesis tested in the computational study presented below.

Model

The goal of the modeling study was to test whether a system encoding grammatical gender can be learned based solely on phonological and semantic information and no explicit representation of gender. The model was also used to test the hypothesis that grammatical gender can be viewed as a higher order regularity in the similarity structure of the training corpus. This was explored by analyzing the extent to which the model's performance in generalization tests is influenced by the coherent co-variation in the semantics<->phonology mapping.

Architecture and Training

The model is a recurrent network with three layers of binary units interconnected using weighted connections (Figure 2). The input to the model was a semantic representation of the word which contained 975 features, 966 of which were semantic features from the norms of McRae and colleagues (e.g., "lives in a nest", "is round", etc.), 2 represented number (singular and plural), and 7 represented case (nominative, genitive, dative, accusative, vocative, instrumental, and locative). Importantly, there was no information about grammatical gender. The output layer consisted of a CCVCC syllabic frame (C=consonant, V=vowel). Each phoneme was coded as a 16-bit vector of phonetic features based on the standard description of Serbian phonemes. The phonological output developed over time, with one syllable output at each time step. This kind of representation permitted the use of multisyllabic words of varying lengths.

The model's task was to produce the word's phonology in the required inflectional form given its semantic representation as input. For example, the word /avion/(airplane)-NOM.SG. was represented on the input layer with the following 15 units on: flies, found in airports, has a propeller, has engines, has wings, crashes, fast, large, made of metal, requires pilots, used for transportation, used for travel, used for passengers, nominative, singular; the rest of the units (960) were off. Activation propagated through the network, and at the output it was trained to produce the phonetic representation for the syllables /a/-/vi/-/on/, one syllable at a time.

The network was trained using the corpus described above. It consisted of 184 nouns (1221 inflectional forms total, 10 items were preserved for the generalization test): 99 masculine (652 inflectional forms) and 85 feminine (569 inflectional forms). The amount of exposure to each word was frequency weighted (Kostić, 1999). The network was trained using the backpropagation through time learning algorithm, with cross-entropy as the error measure. The learning rate was set to 0.005, and the error radius was 0.1. The range of the initial random weights was set to 0.01. The network was run 5 times with different starting random weights, to examine variation in possible network solutions.

Results

Learning The model's performance was assessed using two measures: a coarse-grained measure of the percentage of correctly produced items at the conclusion of training, and a more fine-grained measure, cross-entropy error, which indicates the extent to which the obtained activation on each output unit deviates from the target activation. The items were scored as correct if all its phonemes were closer to the target/correct phonemes than to any other phonemes, using the nearest neighbor criterion based on Euclidean distance. Even when all items in the training corpus are considered learned according to this measure, there is still variation in terms of the distance between the obtained unit activation and the target activation. This variability is reflected in the cross-entropy measure.

The learning curve averaged across five runs is presented in Figure 3. The network starts learning relatively quickly (on average, it acquires almost 20% of the training corpus in the first 40 thousand iterations), and it acquires 85% of the training corpus after 140 thousand iterations. It takes another 400 thousand iterations to acquire the rest of the corpus. Importantly, within one million iterations the network successfully learns the whole training corpus. This can be taken as evidence that learning the inflectional forms of gendered nouns is possible in a system which encodes probabilistic semantic and phonological properties of words *without* an explicit representation of grammatical gender.

In order to test that the model solved this task in a way similar to other models of the same kind, a multiple regression was run with form frequency, lemma frequency (summed frequency of all inflectional forms for a word) and word length (in number of syllables) as predictors and mean cross-entropy error after 1 million iterations as the dependent variable. All predictors accounted for a significant proportion of variance

Figure 3: Learning in the model.

($r^2 = .073$ for word length, $r^2 = .042$ for log lemma frequency, $r^2 = .018$ for log form frequency, all $ps < .001$). This shows that the model's performance was influenced by the factors that are typically relevant for the models of this kind, and also influence human performance (e.g. Harm & Seidenberg, 2004).

Generalization A generalization test consists of presenting a fully trained model with the items not presented during training and recording its output. This test shows the extent to which the model discovered regularities that exist in the training corpus rather than just memorized individual items. We also used this test to explore the extent to which the model's performance was influenced by the structure of the semantics<->phonology mapping as related to grammatical gender.

The generalization corpus consisted of 52 items (10 preserved from the training corpus, and 42 inflectional forms generated from the forms missing in the Frequency dictionary of Serbian (Kostić, 1999)). A subset of 28 items was designed to test the extent to which the model relies on the S<->P consistency, described below. The whole corpus of 52 items was used to assess average generalization through training.

The maximum performance, 68% of the items in the generalization corpus correctly produced, is reached after 360 thousand iterations when on average 99.5% of the items in the training corpus are correctly produced. This performance may be considered satisfactory given the size and representativeness of the training corpus, and the complexity of the Serbian inflectional morphology. For comparison, in a study using a larger corpus (3244 words) the average accuracy on the generalization test was 84.6% (Mirković et al., 2008). It should be noted that it is not clear what is the extent to which human performance in a similar test, having been trained on a similar corpus, would differ from the model. One way to assess this is to train the model on nouns more representative of the early child vocabulary and compare its performance to that of children. It should be emphasized, however, that in the current study the model was used as a tool to explore the view of grammatical gender as coherent covariation in the semantic and phonological space rather than to model acquisition. Thus the factors that influence the model's performance were

the focus of the study rather than the absolute performance on the tests.

A subset of items in the generalization corpus was used to test the effect of S<->P consistency. Specifically, coherent co-variation of semantic properties with morphophonological properties (S<->P consistency) should facilitate generalization in the network, and may be one of the factors that gives rise to what can be termed as grammatical gender.

The test corpus consisted of 28 feminine nouns from 6 semantic categories. Half of the nouns were from semantic categories where all or the majority of items were feminine nouns (buildings, kitchen containers, fruit). In other words, the items in this group coherently co-vary in the semantic and phonological space, so they are considered consistent in the S<->P mapping. The other half of items consisted of nouns from semantic categories where feminine nouns were a minority (birds, bugs, weapons). In other words, the items in this group were similar/coherently co-varied in the semantic space but were "outliers" in terms of morphophonological properties, as the majority of the items in the semantic category had morphophonological properties of masculine nouns. Thus they are considered inconsistent in the S<->P space. The items in the two categories were matched in lemma frequency, length and inflectional form.

The cross-entropy error for the items was submitted to an ANOVA with S<->P consistency as the independent variable, and simulation run as a random factor. The difference between the two conditions was significant, $F(1, 4) = 58.65$, p = .002, such that the cross-entropy error was larger for items from the categories with inconsistent S<->P mapping.

This finding indicates that coherent covariation in the semantic and phonological space improves generalization. Thus the regularities in the words whose properties coherently co-vary in both the semantic and phonological space are easier to discover in learning and to apply in generalization.

General Discussion

In summary, the findings from the corpus analysis suggest that grammatical gender in Serbian nouns is correlated with their morphophonological and semantic properties. Similar properties have been described in other languages (e.g. Corbett, 1991; Kelly, 1992), and human speakers have been shown to encode them (e.g. Brooks, Braine, Catalano, & Brody, 1993; Gerken, Wilson, & Lewis, 2005; Schwichtenberg & Schiller, 2004). A PDP network without an explicit representation of grammatical gender was trained to produce inflectional forms of Serbian nouns based on distributed semantics as input and distributed phonology as output, and no explicit representation of grammatical gender. It successfully learned a training corpus of 1221 inflectional forms. Based on computational principles of PDP models it was suggested that grammatical gender can be viewed as coherent co-variation of features in the semantic and phonological space, and in the S<->P mapping. One implication of this view is that the performance of the model will be facilitated by the consistency

in the semantics<->phonology mapping. This was confirmed in the generalization test.

MacWhinney, Leinbach, Taraban, and McDonald (1989) used a similar type of network to explore grammatical gender processing. However, the input representations in their study were hand-crafted to capture only the information shown to be correlated with grammatical gender. This does not reflect what is available to the child. The model presented here, however, does not include any weighting of either semantic or phonological information in terms of their correlation with grammatical gender; rather, during learning the network discovers the regularities important for successfully solving the task (producing the phonological form based on semantic features). In this sense the model presented here better approximates the task the child learning a language faces.

The findings presented here have implications for models of acquisition and processing of syntactic categories in general. They indicate that syntactic categories may be just a higher order regularity that emerges through learning lower level properties of the linguistic input (see also Elman, 1990; Gasser & Smith, 1998), and so challenge the idea that abstract syntactic features may be necessary for language processing, as suggested by traditional models of language production (e.g. Caramazza, 1997; Levelt et al., 1999).

Acknowledgments

References

Bates, E., Devescovi, A., Hernandez, A., & Pizzamiglio, L. (1996). Gender priming in Italian. *Perception and Psychophysics*, *58*, 992-1004.

Brooks, P. J., Braine, M. D. S., Catalano, L., & Brody, R. E. (1993). Acquisition of gender-like noun subclasses in an artificial language: the contribution of phonological markers to learning. *Journal of Memory and Language*, *32*, 76-95.

Caramazza, A. (1997). How many levels of processing are there in lexical access? *Cognitive Neuropsychology*, *14*, 177-208.

Cassidy, K. W., & Kelly, M. H. (1991). Phonological information for grammatical category assignment. *Journal of Memory and Language*, *30*, 348-369.

Corbett, G. G. (1991). *Gender*. Cambridge, UK: Cambrigde University Press.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179-211.

Gasser, M., & Smith, L. B. (1998). Learning nouns and adjectives: a connectionist account. *Language and Cognitive Processes*, *13*, 269-306.

Gerken, L. A., Wilson, R., & Lewis, W. (2005). Infants can use distributional cues to form syntactic categories. *Journal of Child Language*, *32*, 249-268.

Gonnerman, L. M., Seidenberg, M. S., & Andersen, E. S. (2007). Graded semantic and phonological similarity effects in priming: Evidence for a distributed connectionist approach to morphology. *Journal of Experimental Psychology: General*, *136*, 323-345.

Harm, M. W., & Seidenberg, M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*, *111*, 662-720.

Kelly, M. H. (1992). Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review*, *99*, 349-364.

Kostić, D. (1999). *Frekvencijski rečnik savremenog srpskog jezika [Frequency dictionary of contemporary Serbian language]*. Yugoslavia: Institute for Experimental Phonetics and Speech Pathology and Laboratory for Experimental Psychology, University of Belgrade.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*, 1-75.

MacWhinney, B., Leinbach, J., Taraban, R., & McDonald, J. (1989). Language learning: cues or rules? *Journal of Memory and Language*, *28*, 255-277.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavioral Research Methods*, *37*, 547-559.

Mirković, J., Seidenberg, M. S., & Joanisse, M. F. (2008). *Statistical nature of inflectional structure: Insights from a highly inflected language.* (Manuscript submitted for publication)

Monaghan, P., Chater, N., & Christiansen, M. H. (2005). The differential role of phonological and distributional cues in grammatical categorisation. *Cognition*, *96*, 143-182.

Pinker, S. (1984). *Language learnability and language development*. Cambridge, Massachusetts: Harvard University Press.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: a parallel distributed processing approach*. Cambridge, MA: MIT Press.

Schwichtenberg, B., & Schiller, N. O. (2004). Semantic gender assignment regularities in German. *Brain and Language*, *90*, 326-337.

Shi, R., Morgan, J., & Allopenna, P. (1998). Phonological and acoustic basis for earliest grammatical category assignment: a cross-linguistic perspective. *Journal of Child Language*, *25*, 169-201.

Zubin, D. A., & Köpcke, K. M. (1986). Gender and folk taxonomy: The indexical relation between grammatical and lexical categorization. In C. G. Craig (Ed.), *Noun classes and categorization: Proceedings of a symposium on categorization and noun classification, Eugene, Oregon, October 1983* (p. 139-180). Amsterdam: Benjamins.