

LANGUAGE AND COGNITIVE PROCESSES, 2004, 19 (0), 000-000

Where does gender come from? Evidence from a complex inflectional system

Jelena Mirković, Maryellen C. MacDonald and
Mark S. Seidenberg
University of Wisconsin, Madison, WI, USA

Although inflectional morphology has been the focus of considerable debate in recent years, most research has focused on English, which has a much simpler inflectional system than in many other languages. We have been studying Serbian, which has a complex inflectional system that encodes gender, number, and case. The present study investigated the representation of gender. In standard theories of language production, gender is treated as an abstract syntactic feature segregated from semantic and phonological factors. However, we describe corpus analyses and computational models which indicate that gender is correlated with semantic and phonological information, consistent with other cross-linguistic studies. The research supports the idea that gender representations emerge in the course of learning to map from an intended message to a phonological representation. Implications for models of speech production are discussed.

Language production involves translating a conceptual representation, a message, into a phonological code that serves as the basis for an articulatory plan. Within this broad framework, there have been several proposals concerning the nature of the representations that mediate this

Correspondence should be addressed to Mark S. Seidenberg, Department of Psychology, University of Wisconsin, Madison WI 53706. Email: marks@lcnl.wisc.edu.

This research was supported by NIMH grant RO1 MH58723 and NIMH grant P50 MH64445. Mark Seidenberg was also supported by an NIMH research scientist development award. We would like to thank Marc Joanisse with whom we collaborated on developing the models used here, using software developed by Michael Harm, whom we also thank. We also thank Aleksandar Kostic for making available the electronic version of the Frequency Dictionary of Contemporary Serbian. This work benefited from the feedback of the attendees of the 3rd Workshop on Morphological Processing, held in Aix-en-Provence, France.

© 2004 Psychology Press Ltd

<http://www.tandf.co.uk/journals/pp/01690965.html> DOI: 10.1080/01690960444000205

2 MIRKOVIĆ ET AL.

computation and the degree of interactivity among different types of representations. A complicating issue is the fact that words also encode information about the syntactic structures in which they participate, information that also influences a word's phonological form. Many languages employ inflectional affixes that convey largely grammatical information such as number, tense, case, and gender. Whereas English nouns encode only number inflectionally (e.g., *cat* - *cats*), many other Indo-European languages require that nouns (and often also adjectives and determiners) mark grammatical gender (as in the Spanish *el chico* (boy, masculine) vs. *la chica* (girl, feminine)). Some of these languages also mark case, indicating the grammatical role of a noun phrase in a sentence. One such language is Serbian, in which nouns are marked for case, number and gender.¹ For example, the suffix /u/ in /kravu/[cow] denotes that the noun is feminine accusative singular.²

Models of word production have traditionally assumed that gender and other grammatical features are represented independently from a word's semantics and phonological form. For example, although the word production models of Levelt and colleagues (Levelt, Roelofs, & Meyer, 1999) and Caramazza (1997) differ in some respects, they share two important assumptions. First, syntactic information, including grammatical gender, is represented by abstract features at a distinct syntactic level of representation. Second, processing at this syntactic level is independent of phonological and semantic processing (see Figure 1).

Evidence for the claim that gender is represented as an abstract syntactic feature has been obtained using two language production tasks, picture-word interference and fragment completion. In the former task (e.g. Glaser & Dungelhoff, 1984; Lupker, 1979; Schriefers, Meyer, & Levelt, 1990), the participant names a picture using a single word or, in some studies, a noun phrase (e.g. "dog", "a dog", "big dog"). A distractor word is presented on or near the picture (in case of auditory distractors, in close proximity or simultaneously with presentation of the picture). Depending on the timing of distractor presentation and other properties of the materials, the distractor word may interfere with picture naming, yielding longer naming latencies. Gender has been studied using this task by manipulating the genders of the pictured item and the distractor word. Several studies in Dutch found that picture naming latencies were longer when the distractor word had a different gender than the pictured item, compared with when they had the same gender (e.g., Schriefers, 1993).

¹ Languages referred to as Serbian, Croatian, and Bosnian were referred to as Serbo-Croatian in earlier literature.

² All Serbian examples are written in International Phonetic Alphabet, unless otherwise specified.

REPRESENTATION OF GENDER 3

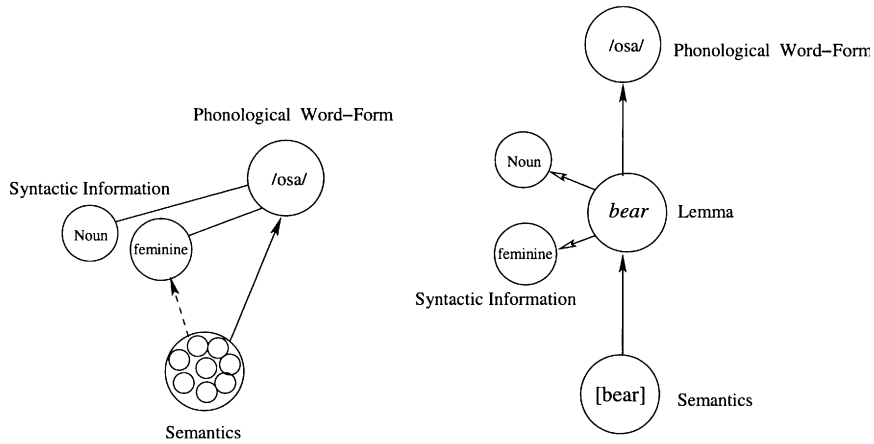


Figure 1. Word production models showing activation of the Spanish word /osa/[she-bear]. The model based on Caramazza (1997) is presented on the left, and the model based on Levitt et al. (1999) is presented on the right. Both models represent grammatical gender as an abstract syntactic node.

This result has been taken as evidence for priming of an abstract gender node, in that the gender of the distractor word interferes with activating the appropriate gender of the pictured item.

The fragment completion paradigm has been used to investigate agreement processes, whereby grammatical features on a noun (gender, number, and/or case) are also marked on other words in the sentence, as when a determiner or adjective agrees with the noun it modifies. In this paradigm (e.g. Bock, Eberhard, Cutting, Meyer, & Schriefers, 2001; Vigliocco, Hartsuiker, Jarema, & Kolk, 1996), participants are given a short sentence fragment and must generate a sentence that begins with this phrase. The dependent variable is the rate of agreement errors produced in the sentences. One focus of this research is whether semantic or phonological properties of the agreeing words influence the agreement process. In several studies these properties did not modulate the rate of agreement errors (e.g. Bock & Eberhard, 1993). These results have been taken as support for the view that agreement is an autonomous process in which abstract grammatical features such as gender or number are copied from a noun to the agreeing forms.

ALTERNATIVE VIEWS OF GENDER REPRESENTATION

Although these data are consistent with the view that word production involves activation of abstract gender and other grammatical features,

4 MIRKOVIĆ ET AL.

other data raise questions about this approach. First, there is considerable linguistic research emphasising the extent to which grammatical gender is related to semantics and phonology (e.g., Corbett, 1991). In a comprehensive study of gender, Corbett (1991) argued that grammatical gender is predictable from various semantic and morphophonological regularities. He described the semantic and phonological correlates of gender across languages, attempted to describe them in terms of rules, and concluded that there may not be a need for a syntactic specification of gender. Corbett's view is that all gender systems have a semantic core; languages vary in the extent to which they make use of additional morphophonological information. In French, for example, a language whose gender system is often thought to be highly arbitrary, 94.2% of nouns ending in /ʒ/ are masculine, whereas 90% of nouns ending in /z/ are feminine (Corbett, 1991). Corbett (1991) also suggested that these language-specific correlates of gender are learned through exposure to a broad range of utterances.

Corbett (1991) offered similar observations about semantic regularities correlated with gender. He argued that gender has a semantic basis but not always the male–female distinction that is typical of Indo-European languages. Many languages, such as Ojibwa and other Algonquian languages, use an animate-inanimate distinction as a basis for the gender system. Moreover, even in Indo-European languages such as German, which incorporates clear male–female distinctions in the gender system, there are groups of same-gendered nouns that cluster on the basis of other semantic properties (e.g., Zubin & Köpcke, 1981, 1986). German nouns denoting superordinate categories are usually neuter, for example *Obst* (fruit), *Instrument* (instrument), *Gemüse* (vegetables); nouns denoting alcoholic drinks tend to be masculine; nouns denoting apes, mammals, and birds tend to be masculine, whereas nouns denoting reptiles and lower animals tend to be feminine. Corpus analyses supported these generalisations: for example, Zubin and Köpcke (1986) estimated that 70% of nouns denoting birds are masculine, whereas 76.9% of the nouns for lower animals are feminine (Zubin & Köpcke, 1986). Corbett (1991) identified gender systems based on semantic distinctions including animate–inanimate, natural gender, whether the word refers to something rational or non-rational, human vs. non-human, and augmentative vs. diminutive. These criteria vary in frequency across languages, with animacy and natural gender the most common.

These patterns appear to be part of a broader trend in which syntactic elements turn out to have semantic and/or phonological correlates. Grammatical categories such as noun and verb provide a salient example (e.g. Miller & Fellbaum, 1991; Kelly, 1992). The semantic correlates of these categories have been widely recognised (e.g., Gentner, 1981; Miller & Fellbaum, 1991); the phonological regularities associated with these

REPRESENTATION OF GENDER 5

classes have been investigated in recent years because of their roles in acquisition and processing. Kelly (1992) showed that phonological characteristics of English nouns and verbs vary in terms of factors such as stress, number of syllables, and word duration and that people make use of this information in processing. Similarly, Haskell, MacDonald, and Seidenberg (2003) identified phonological regularities among adjectives in English. These regularities were discovered by a simple connectionist model trained to distinguish adjectives from other open class words on the basis of phonological information alone. The model correctly classified 80% of the words in the training set. Cassidy, Kelly, and Sharoni (1999) showed that even in English, which codes grammatical gender only in pronouns, male and female names tend to have different phonological properties, which people use to infer the natural gender of the named individual (see also Cutler, McQueen, & Robinson, 1990). Studies of language acquisition suggest that these phonological cues play an important role in the acquisition of grammatical categories (Gerken, 2002). Thus the linguistic analyses indicate that there are strong correlations between grammatical features such as gender and semantic and phonological properties of words; the studies of acquisition and processing suggest that this information is utilised.

Some studies of language production also call into question the idea that gender is represented by a feature at an autonomous level of syntactic representation. Caramazza and colleagues observed that the gender interference effect in picture-word interference tasks, which has been taken as evidence for abstract gender nodes (e.g. Schriefers, 1993), does not occur in all languages. In Romance languages the form of a gender-marked determiner depends not only on the gender of the noun but also may depend on the phonological form of the noun (as in English, where the phonological form of the noun determines the choice between the *a* and *an* determiners). Caramazza, Miozzo, Costa, Schiller, and Alario (2001) interpreted this result to indicate that the architecture of the production system must afford different degrees of interactivity between different types of representations depending on the language. Thus their syntax and phonology interact in Romance languages but not in Germanic languages, in which phonological form does not affect determiner selection. However, these results could equally be taken as evidence against autonomous gender nodes, in favour of a representation that integrates semantic and phonological information. Languages apparently differ in the extent to which gender relies on each of these factors. The same issue arises in connection with several studies using the fragment completion paradigm suggesting that both broad semantic factors beyond the conceptual representation of individual nouns (e.g., Thornton & MacDonald, 2003; Vigliocco et al., 1996) and morphophonological factors

6 MIRKOVIĆ ET AL.

(e.g., Haskell & MacDonald, 2003; Vigliocco, Butterworth, & Semenza, 1995) influence the process of subject–verb number agreement.

In summary, considerable progress has been made within theories in which gender is treated as an abstract syntactic feature. At the same time, there is substantial evidence that gender is correlated with semantic and phonological factors that are segregated from syntactic representations in such theories, and that this information is used in acquisition and processing. Minimally these findings raise questions about how these correlations can be represented in models such as those illustrated in Figure 1. A further possibility is that the use of gender information in processes such as agreement is not merely influenced by semantic and phonological information; rather, that is *how gender is represented*. Gender, on this view, is an emergent property in lexical systems in which statistical relations between different types of information (principally semantics and phonology) within and across words are encoded. Note that on this view no single cue has to be entirely reliable, and indeed cross-linguistic analyses suggest that they are not. A given cue does not have to occur in all or even a majority of cases to be used in networks that encode probabilistic constraints (e.g., connectionist constraint satisfaction networks; Seidenberg & MacDonald, 1999). Such networks can make use of even lower probability constraints; moreover, the interactions among constraints are nonlinear. Thus, two cues that are only mildly constraining in isolation may be highly constraining when taken together. Hence the existence of substantial though nonetheless imperfect cues to gender such as the ones discussed by Corbett (1991) and Zubin and Köpcke (1981, 1986) are significant.

In this article we investigate the role of semantic and phonological factors in gender representation and word production. We first present corpus analyses of gender in Serbian, a language with a complex inflectional system, including gender marking. These data indicate that semantic properties are highly but not perfectly correlated with gender in this language. Thus semantics provides strong probabilistic cues to gender. We then describe a computational model of noun production used to investigate the gender system. This model was developed by Mirković, Seidenberg, & Joanisse (2004) as a preliminary exploration of the capacity of connectionist networks to learn a complex inflectional system. The model (described further below) performed a simplified version of a production task: it was given information about a word (e.g., lemma, number, case, and sometimes gender) as input and had to produce its correct phonological form as output. The model succeeded in learning a large corpus of inflected words, generalised on the basis of this knowledge, and provided information about sources of complexity in the Serbian system. In the present research, we examined gender issues by comparing

REPRESENTATION OF GENDER 7

the performance of three versions of this basic model. In one condition the model included an explicit representation of a noun’s gender as part of the input. In a second condition, the model did not receive any information about gender; it could therefore only produce correctly inflected forms by picking up on other types of information that are correlated with gender. In a third condition the model included a semantic cue that is strongly correlated with gender in Serbian. The comparisons between these conditions provide evidence concerning the extent to which gender marking can arise from correlated information.

SEMANTIC CORRELATES OF GENDER IN SERBIAN NOUNS

Serbian is a south Slavic language with a complex inflectional system (see Table 1). Nouns are coded for case (7 of them), number (singular-plural) and gender (masculine, feminine, and neuter). As the examples in Table 1 illustrate, because Serbian is a fusional language, these properties are coded with a single suffix. However, the relationship between the suffixes and the three properties is probabilistic. For example, with respect to gender, most masculine nouns in nominative singular end in a

TABLE 1
Inflectional forms of nouns /medved/ [bear], /krava/ [cow] and /selo/ [village]

case	Singular forms		
	masculine	feminine	neuter
nominative	medved	krava	selo
genitive	medveda	krave	sela
dative	medvedu	kravi	selu
accusative	medveda	kravu	selo
instrumental	medvedom	kravom	selom
locative	medvedu	kravi	selu
vocative	medvede	kravo	selo

case	Plural forms		
	masculine	feminine	neuter
nominative	medvedi	krave	sela
genitive	medveda	krava	sela
dative	medvedima	kravama	selima
accusative	medvede	krave	sela
instrumental	medvedima	kravama	selima
locative	medvedima	kravama	selima
vocative	medvedi	krave	sela

Note: All Serbian examples are written in the International Phonetic Alphabet.

8 MIRKOVIĆ ET AL.

TABLE 2
Examples of singular forms of a masculine (/sudija/[judge]) and a feminine noun /domatçitsa/[housewife]) ending in /a/

<i>case</i>	<i>masculine</i>	<i>feminine</i>
nominative	sudija	domatçitsa
genitive	sudije	domatçitse
dative	sudiji	domatçitsi
accusative	sudiju	domatçitsu
instrumental	sudjom	domatçitsom
locative	sudiji	domatçitsi

consonant. However, there is a small group of masculine nouns that end in /a/ in this case/number form (e.g., /sudija/[judge], Table 2). Ambiguity is created by the fact that the /a/ ending is also used for the majority of feminine nouns in nominative singular. Interestingly, however, these masculine nouns are semantically similar: most of them refer to professions or actions that were traditionally performed by men, such as /sudija/[judge], /vodza / [leader]. To complete the circle, some feminine nouns end in a consonant in this case/number form (e.g., /strast/[passion], Table 3).

As Corbett (1991) noted, in some languages a more reliable correlation between gender and form is obtained by looking at the whole set of inflectional forms a noun can take. That is, the gender of nouns ending in a consonant cannot be determined by looking only at the nominative singular (citation) form; however, a clear distinction between masculine and feminine nouns is observed if all inflectional forms are taken into account (Table 3). Because all inflectional forms of a noun have to be taken into consideration in order to find a reliable correspondence to grammatical gender, Corbett classifies the Serbian gender system as morphological (Corbett, 1988, 1991).

TABLE 3
Examples of singular forms of a masculine (/rast/[growth]) and a feminine noun (/strast/[passion]) ending in a consonant. Note the difference in the suffixes of the nouns of the two genders

<i>case</i>	<i>masculine</i>	<i>feminine</i>
nominative	rast	strast
genitive	rasta	strasti
dative	rastu	strasti
accusative	rast	strast
instrumental	rastom	strastju
locative	rastu	strasti

In Serbian and other Slavic languages, gender systems are based on a male–female distinction for humans and for animals whose sex matters to humans (Corbett, 1988). For example, /lav/[lion] is masculine, whereas /lavitsa/[lioness] is feminine. Many neuter nouns refer to human or animal offspring, e.g., /devojtse/[young girl], /kutse/[puppy], /matse/[kitty], /lane/[fawn]. Consistent with Zubin and Köpcke’s results for German (Zubin & Köpcke, 1981, 1986), these examples suggest there may be strong semantic correlates of grammatical gender in Serbian nouns.

We investigated this possibility with respect to two potential semantic cues, animacy and abstractness, which were chosen because they are correlated with gender in other Indo-European languages (Corbett, 1991). We extracted a sample of 407 Serbian noun lemmas from the Frequency Dictionary of Contemporary Serbian (Kostić, 1999) to identify the distribution of the two possible gender cues. Approximately 74% of these nouns (300 lemmas) were randomly drawn from the Serbian frequency dictionary, and the rest were added to represent nouns that have relatively less frequent morphophonological patterns. The added items (107 lemmas) were randomly selected from the nouns with these lower frequency patterns. This corpus was also used in the simulation work presented below, and more details about the sample are provided there. Preliminary simulations run without these added items indicated that adding them had little impact on the model’s learning of the training set; however, they do improve generalisation because the model cannot produce generalisations of lower frequency patterns unless it has been exposed to them. Each noun was coded for the two semantic features. Animacy was coded by the first author as a binary factor. Abstractness-concreteness is a graded dimension, but for the purpose of this work, we coded it in binary terms. To assign these values, we used English translations for Serbian nouns and the corresponding concreteness norms in the MRC Database (Coltheart, 1981). All nouns that had a concreteness rating higher than the mean concreteness rating in the database were coded as concrete, and the ones below were coded as abstract. In cases where the direct translation was not in the database, the closest semantic neighbour was used (e.g., *pound* was used instead of *kilo*).

The data presented in Figure 2 show that animacy, and to a lesser degree, abstractness, are strong cues to gender in Serbian. The vast majority of animate nouns in the sample (81.08%) are masculine, and the remaining are equally divided between feminine and neuter. Similarly, over half of the abstract nouns are feminine.

Given these positive results for abstractness and animacy, we next investigated a narrower semantic distinction. In some languages, food-related properties of words are taken as one of the criteria for gender classification (Corbett, 1991). In many languages, nouns that refer to edible

10 MIRKOVIĆ ET AL.

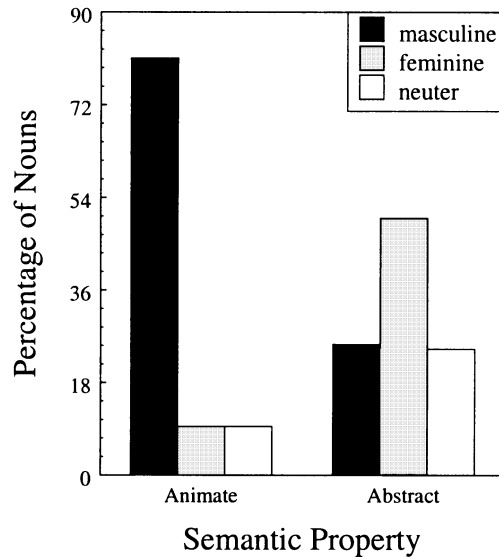


Figure 2. Animacy and abstractness as cues for gender in Serbian.

and non-edible things belong to different genders (Corbett,1991). In Dyirbal, an Australian language, an even narrower semantic distinction is informative: Dyirbal has a gender for non-flesh food. We investigated whether something similar occurred in Serbian, specifically whether the categories of fruit vs. vegetable co-vary with gender. The sample of nouns used above contained only one noun from these categories and so could not be used for this analysis. We therefore used Serbian translations of the English words that were unambiguously classified as a fruit or as a vegetable in the production norms of McRae, Cree, Seidenberg, and McNorgan (2004), which yielded 18 fruits and 20 vegetables. The data presented in Figure 3 show that nouns referring to fruits in Serbian tend to be feminine (72.22%) and nouns referring to vegetables tend to be masculine (65%). Thus at least some narrower semantic distinctions also co-vary with gender in Serbian.

In summary, these data suggest that, as in other languages, the Serbian gender system is at least partially determined by the morphophonological and semantic properties of nouns. Of course the cues we examined are probabilistic rather than absolute; given these semantic properties one cannot predict the gender of a word with certainty. However, we examined only two cues; the interesting possibility is that the system as a whole entails multiple probabilistic cues that together establish a word's gender with a high degree of certainty. This is a basic characteristic of constraint satisfaction processing systems. Cues that may be only partially reliable

REPRESENTATION OF GENDER 11

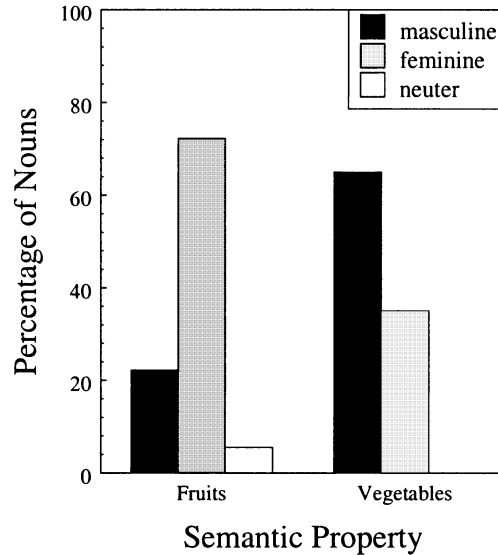


Figure 3. Fruits and vegetables as cues for gender in Serbian.

when taken individually become highly constraining when taken in conjunction with other cues.

COMPUTATIONAL EVIDENCE

We then wanted to explore whether the representation of gender in terms of semantic and phonological regularities could provide a basis for the representation of gender in models of word production. As a tool, we used a connectionist (PDP) model that was developed to address morphological processing in Serbian nouns (Mirković et al., 2004). The earlier work had two main goals. One was to determine, in a preliminary way, whether a complex inflectional system such as in Serbian could merely be encoded by a simple connectionist network. Traditional grammars of the language describe the system in terms of a large number of complex rules with many conditions attached to them. Such a system might better be described by a set of simultaneous probabilistic constraints, and we built a model to examine this. The other was to use the model as a tool for discovering sources of consistency and complexity in the system that are otherwise difficult to identify. The model took an explicit representation of lemma, gender, number, and case as input and produced the correctly inflected phonological form as output. Details are provided in the Mirković et al. (2004) article and below; the model was based on an architecture developed by Joanisse (2000) to study Dutch inflectional morphology. The model learned to produce correctly inflected forms for a large corpus

12 MIRKOVIĆ ET AL.

of Serbian nouns (over 3000 words). The model revealed some additional non-obvious facts about the system. One was that, in the model, the main source of complexity was learning about deformations of the stem (e.g., epenthesis, palatalisation) that are influenced by the inflection, rather than learning the inflections themselves. For example, if the phoneme preceding the suffixes /e/ or /i/ is /k/, /g/, or /x/, the consonant changes to /ts/, /z/, or /s/, respectively, in case of suffix /i/ and to /tʃ/, /ʒ/, /ʃ/ respectively in case of suffix /e/, as in /savetnik/[counselor]-MASC.NOM.SG., /savetnitsi/-NOM.PL, /savetnitʃe/-VOC.SG. The model also led to the discovery of a novel type of lexical neighbourhood that was shown to influence reading performance in native speakers of Serbian. In summary, this research represented a first attempt to model a complex inflectional system using principles that had been previously explored in work on the simpler English inflectional system and the representation of spelling-sound information.

The present research built on this work but focused on a different question: whether gender emerges from semantic and phonological regularities. The main data concern the performance of three versions of the Mirković et al. (2004) model that differed with respect to the way grammatical gender was represented.

Description of the Models

The Mirković et al. (2004) model performed a simplified word production task (Figure 4, top). In the present study, three versions of the model were compared. One was the model described by Mirković et al. (2004), in which the gender of nouns was explicitly represented by three binary nodes, one for each gender (*Explicit Gender* condition). In the *Morphophonology Only* condition, gender was not represented on the input at all. In the *Animacy* condition, the model again did not have an explicit representation of gender; rather, the input included information about a probabilistic semantic cue, animacy. The principal research question concerned the extent to which a model of this type could produce correct output without an explicit gender feature on the input, using correlated semantic and phonological cues. It was beyond the scope of this research to develop full semantic representations for all words in the training corpus, and so the input did not represent all of the semantic features that may be correlated with gender. Instead we included only animacy, a highly salient feature identified in the corpus analysis described above. Given this limitation, we expected the model to perform worse than the condition in which gender is explicitly represented. However, this could be for either of two reasons. One is because gender really is represented in terms of an abstract syntactic feature, which the model

REPRESENTATION OF GENDER 13

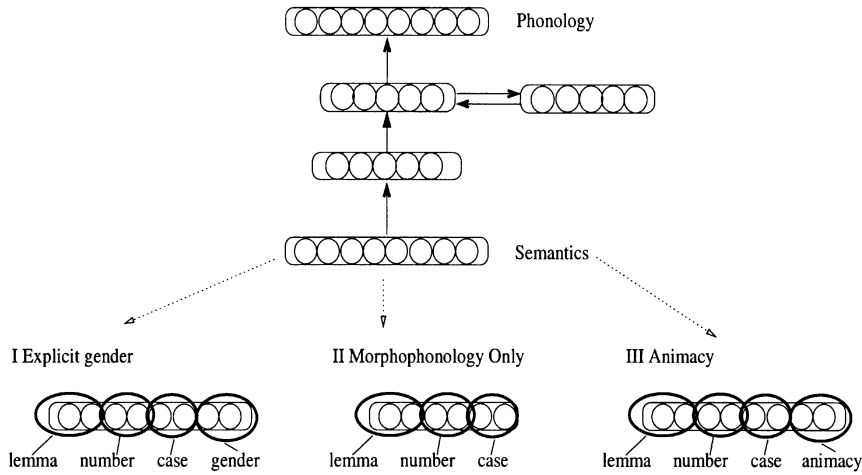


Figure 4. Architecture of the models. The general architecture relevant to all three models is at the top. The bottom row shows the semantic layer in the three types of models. Ellipses represent groups of units organised into layers. Arrows indicate directional weighted connections used to pass information among layers. The semantic layer is connected to the phonological output through two hidden layers with 250 and 100 nodes respectively. The second hidden layer has a layer of recurrence with 100 nodes.

lacked. The other is that the implemented model had limited semantic information with which to work: lemmas were encoded by localist nodes rather than distributed representations of their meanings, and the semantic condition included only one feature, animacy. Thus, the models with no explicit gender information had a built-in limitation on the extent to which correlates of gender could be discovered. Explicit gender nodes could be seen as standing in for these missing semantic representations, something that will need to be addressed in future research by developing models that use richer semantic representations. In this modelling we were mainly interested in examining the extent to which the model could pick up gender regularities given the information it had to work with, not whether performance would match that in the explicit gender condition.

Architecture of the Models

The basic architecture of the models is presented in Figure 4. All models were fully recurrent networks that consisted of four layers of units with weighted connections between them. Each model’s task was to produce the phonological form of a word given different types of input information described below. We refer to this input layer as “semantic” because in an ideal model it would consist of a distributed representation of the intended meaning (see Harm & Seidenberg, in press, for a model using this type of

representation). We have assigned specific functions to units for computational simplicity and labelled them for clarity. In reality we think these representations are also learned.

The output layer represents the word's phonology as a sequence of syllables that develops over time. The syllables are vowel centred within a CCVCC frame (C = consonant, V = vowel; /r/, /l/ and /n/ were coded as vowels in cases where they had a vocalic function; Stanojčić, Popović & Micič, 1989). Each phoneme in the syllable was represented by 16 binary phonetic features (Table 4), based on the standard description of Serbo-Croatian phonemes (Stanojčić et al., 1989). The use of recurrence allowed the network to generate a series of discrete outputs over time, which permitted simulating the production of multisyllabic words of varying lengths. Each syllable is produced on a different time step.

In all of the models, there were localist units representing a word's lemma, number and case. That yielded 407 lemma units, 2 number units (singular, plural), and 7 case nodes. Localist representations were used instead of distributed ones in order to keep the simulation computationally tractable. (Even in this simplified form, each model took several days to run.) To be clear, the use of localist representations here was a pragmatic choice, not a theoretical claim. In future research it will be important to determine whether, as we assume, the results replicate using distributed representations (e.g., semantic feature representations of the meaning associated with a given lemma rather than a lemma node). Ideally the input layer should reflect as closely as possible the nature of the production task, which involves starting with an intended message, not explicit representations of lemma, gender, number, and case.

The three conditions are illustrated in Figure 4. The *Explicit Gender* condition included units encoding lemma, number, case, and also gender. In the *Morphophonology Only* condition, the input layer consisted of units representing lemma, number, and case, but not gender. Because gender is not explicitly represented and the semantics of the lemma was represented by localist units, the model could only derive the correct inflectional form based on morphophonological information. The *Animacy* model was identical to the *Morphophonology Only* condition except that one node that represented the animacy of the noun was added to the input layer.

Training

The training procedure was the same in all three conditions. The input was the representation of a word, which varied across conditions. For example, the model was presented with the word /krava/[cow] either as <krava>-NOM.SG.FEM. in the *Explicit Gender* condition, or as <krava>-NOM.SG. in the *Morphophonology Only* condition, or as

<krava>-NOM.SG.ANIMATE in the *Animacy* condition. The model's task was to produce the correct phonological form of the word as output (in this example, the features representing the five phonemes in /krava/). On a training trial the representation of a word was activated on the input layer for a pre-set number of time steps. Activation propagated from the input to other layers. The first syllable was activated on the output layer starting from time step 3. For example, the features coding the phonemes in the syllable /kra/ would be active for time steps 3 and 4, and the features coding the syllable /va/ would be active for time steps 5 and 6. The models were trained using the backpropagation through time learning algorithm (Pearlmutter, 1995) which compares target activation values in the output layer to the obtained activations at each time step, and adjusts each connection weight in a way that gradually minimises differences between obtained and desired activations. Activation was compared using cross-entropy error (Bishop, 1995), which improves learning using backpropagation learning algorithm in periods of stagnation. The error radius was set to 0.1, meaning that errors less than this value were counted as zero, and the learning rate was 0.005. The network was initialised with small random weights.

Training corpus

The training corpus consisted of 3274 nouns generated from 407 lemmas taken from the Frequency Dictionary of Contemporary Serbian (Kostić, 1999). Each noun lemma in Serbian can have between 1 and 14 surface forms (7 cases × 2 numbers); for most lemmas, however, the model was exposed to only a subset of these forms. This reflects the fact that not all inflectional forms appear in the corpus of 2 million words on which the Serbian frequency dictionary was based. Also, we withheld 60 items (one inflectional form for 60 lemmas) to use in the generalisation test discussed below.

Several criteria were applied in the selection of words for the training corpus. The maximum length of words was six syllables and each syllable could not exceed the CCVCC frame. This did not require excluding many items; approximately 0.5% of the nouns in the Serbian frequency dictionary are more than 6 syllables long. All the words were in the Ekavian dialect. The final training corpus consisted of 42.26% masculine nouns, 36.85% feminine nouns, and 20.88% neuter nouns. These proportions are consistent with the proportions obtained in a larger corpus³ (corpus: masculine–44.99%, feminine–40.43%, neuter–14.57% (Kostić, personal communication).

³ This larger corpus consists of 2 million words from the 20th century daily press and poetry (Kostić, 1999, and <http://www.serbian-corpus.edu.yu>).

Word frequencies taken from the Serbian frequency dictionary were logarithmically transformed based on the equation

$$p = \frac{\log(f + 1)}{\log(f_{max})}$$

where f is the frequency of the word, and f_{max} is the frequency of the most frequent word in the corpus (/broj/(number)-MASC.NOM.SG., frequency of 1426 occurrences out of 2 million words). This type of frequency compression preserved the general statistical structure of Serbian, while assuring that the network was exposed to low frequency words without extensive sampling.

RESULTS

Learning

Five replications, corresponding to different “subjects”, were run in each condition. The replications differed in terms of the initial random weight assignments, and because words were selected for training on a quasirandom basis (biased only by frequency), each model was trained on a different sample of words. All the models were trained for 3 million iterations (trials). Data were averaged across simulations, which yielded very similar results.

Performance of the model was assessed by a Euclidean distance nearest neighbour criterion, whereby an item was scored as correct if each computed phoneme was closer to the correct one than to any other one. It took the models approximately one million iterations to learn to produce 90% of the training corpus correctly (Figure 5). Importantly, the three models exhibited different learning curves. The Animacy model correctly produced an average of 92.35% of the items in the training corpus after 900,000 iterations, whereas it took an additional 100,000 iterations for the Morphophonology Only model to reach the same level of performance (as compared to the Explicit Gender model which took 800,000 iterations). For each run of the models a logistic function was fit to the learning curve, which in all cases yielded a good fit, with r^2 in the range 0.80–0.91, and $p < .001$. This procedure yielded estimates of the slope coefficient for each curve. The analysis of variance on the slope coefficients showed a significant main effect of model, $F(2, 12) = 18.781$, $p < .001$, with the Morphophonology Only model having the shallowest slope. Planned comparisons showed that the Animacy model had a steeper learning curve than the Morphophonology Only model, $t(12) = 3.527$, $p < .01$. This analysis shows that the Morphophonology Only model was slower than the Animacy model in learning the training corpus. Importantly, however, all

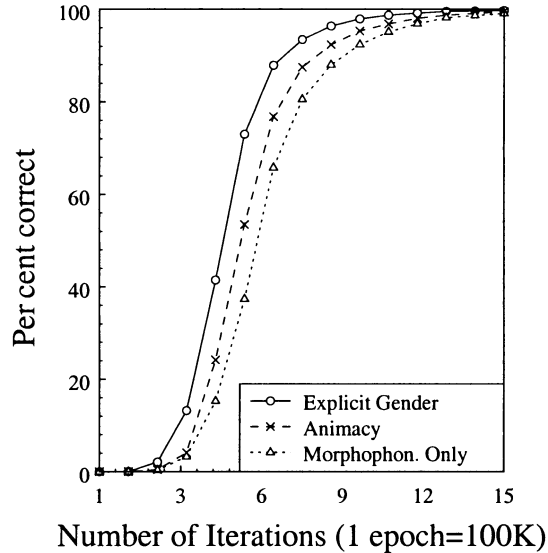


Figure 5. Per cent correct during learning (averaged across 5 runs of each of the models). The first 1.5 million trials are shown in the figure, the point in which all models correctly produced at least 99% of the training corpus correctly.

the models reached essentially perfect performance within 3 million iterations: the Morphophonology Only and the Animacy models produced on average 0.6 items incorrectly, whereas the Explicit Gender model produced no errors.

The items that were incorrectly produced across the five runs of the simulation in the Animacy and the Morphophonology Only models after 3 million iterations were very similar. The errors on the two items which both models missed were identical: /plemitɕ/[nobleman]-MASC.NOM.SG. produced as /plemits/, and /maStɕu/[lard]-FEM.INST.SG. produced as /mastɕu/. In the few items that were not learned correctly the errors were due to missing one or two features in a phoneme.

The results indicate that all the models, regardless of how gender was treated, successfully learned the training corpus, but at different rates. In order to examine this more closely, we performed an analysis of variance on the percentage of errors the models produced after 1.5 million iterations. The analysis of variance on the percentage of errors per gender showed significant main effects of model, $F(2, 12) = 31.591, p < .001$, and gender, $F(2, 24) = 39.745, p < .001$. Importantly, the model x gender interaction was also significant, $F(4, 24) = 3.886, p < .05$. In addition to being correlated with grammatical gender, animacy is relevant for determining the suffix of the accusative singular form of masculine nouns: in animate nouns the accusative singular form is the same as the genitive

singular form whereas in inanimate nouns it is the same as the nominative singular form (e.g., /medved/[bear]-NOM.SG., /medveda/[bear]-ACC.SG. = GEN.SG.; /krompir/[potato]-NOM.SG. = ACC.SG., /krompira/[potato]-GEN.SG.). Thus, the difference between the models might have been carried by the accusative singular forms of masculine nouns. We therefore performed an additional analysis with the accusative singular forms of masculine nouns removed. The analysis showed the same effects: main effects of model, $F(2, 12) = 33.290, p < 0.001$, and gender, $F(2, 24) = 37.356, p < 0.001$, and a model \times gender interaction, $F(4, 24) = 3.698, p < 0.05$. The significant interaction occurred because masculine nouns benefited most from the addition of the semantic cue, as expected (Figure 6).

The main effect of gender can be explained by the statistical structure of phonological properties of nouns of different genders. The largest number of errors in all three models is in masculine nouns and the majority of these (between 62% and 80%) are in the final phoneme of the nominative singular form. Whereas in the majority of feminine and neuter nouns the nominative singular ending is typical for that gender (/a/ and /o/ respectively), any consonant can appear in the nominative singular form of masculine nouns. For example, of the 25 consonants in Serbian only /dʒ/ does not appear as the nominative singular ending in masculine nouns in the Serbian frequency dictionary. Therefore, in terms of type frequency

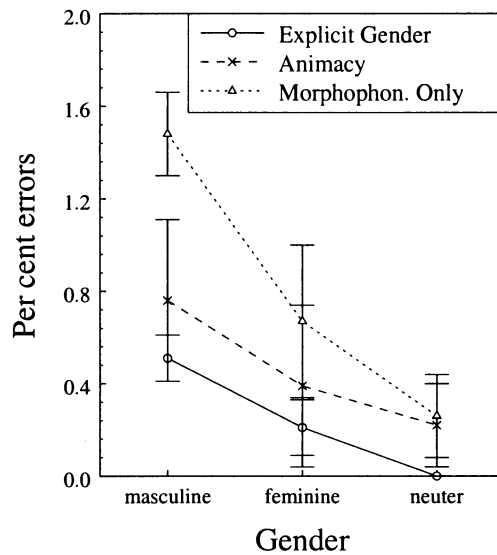


Figure 6. Percentage of incorrectly produced items in the three models after 1.5 million iterations.

20 MIRKOVIĆ ET AL.

any particular ending in the nominative singular of masculine nouns will be less frequent than the one for feminine or neuter nouns.

In summary, the analyses to this point show that grammatical gender in Serbian nouns is largely learnable from morphophonological and semantic regularities that exist in a substantial subset of words. Importantly, the simulations show that morphophonological cues are highly reliable in determining the correct inflectional forms of nouns: the Morphophonology Only model successfully learned more than 99% of the training corpus. However, if the performance of the model is more carefully examined, we can see that the model makes use of the additional semantic cue that was provided in the Animacy model, which speeds up learning.

One question raised by these results is whether the models achieved high levels of performance by merely memorising the training sets. With sufficient training the models could learn input-output mappings for all words without necessarily relying on statistical regularities that occur across words. In order to examine this possibility we conducted further tests of the models' capacities to generalise. Models that have merely memorised the training set would not be able to produce correctly inflected forms for items on which they had not been trained.

Generalisation

The generalisation stimuli consisted of 60 items. These words were derived from the lemmas the models were trained on (20 items per gender), but they were withheld from training. For example, the model was trained on all forms of the word /biʎka/[plant] except for the dative singular form, which was then used in the generalisation test.

The items in each gender were chosen to represent two broad sets: one consisted of nouns that are morphophonologically more typical for one of the genders in terms of type frequency, and the other consisted of nouns that are less typical of that particular gender. For example, approximately 70% of Serbian feminine nouns end in /a/ in nominative singular and their inflectional forms are made by changing the suffix, leaving the stem unchanged (Mirković et al., 2004). By contrast, approximately 10% of feminine nouns end in a consonant in nominative singular (cf. Table 3) and take a different set of suffixes than the rest of feminine nouns. This factor, referred to as category size, was shown to have an effect both in the performance of the model and in word naming in Serbian native speakers (Mirković et al., 2004), such that speakers were slower to name nouns from smaller categories compared with items from larger categories (with words equated in frequency and length). Similarly, the model performed worse on these items, both in learning and generalisation. Therefore the items in the generalisation test were matched across genders for this factor, such

that half of the items in each gender were from the largest category and half were from smaller categories.

Testing involved presenting a generalisation item as the input and determining whether it produced the correct phonological form. Generalisation was tested after every 100,000 learning trials (Figure 7). The Animacy model was always slightly better than the Morphophonology Only and slightly worse than the Explicit Gender model. At asymptote, the Animacy model produced 72.33% of the items correctly, whereas the Morphophonology Only model yielded 68.33% correct and the Explicit Gender model 78%. This performance indicates that the models had not simply memorised the training set. The absolute level of performance is difficult to assess because the model's training set was much smaller than a native speaker's vocabulary. The training corpus included only a subset of the nouns in the language. Perhaps more importantly, it did not include words from other grammatical categories, which is relevant because some morphophonological patterns such as jotation occur more frequently in other categories (e.g., adjectives) than in nouns. Finally, as we have noted, the localist representations of the input did not allow the model to pick up on other correlations between form and meaning. Thus, it is relevant that the models were able to correctly produce many untrained forms using the restricted information in the training corpus, and there is a suggestion that higher levels could be achieved by using richer semantic representations and larger corpus.

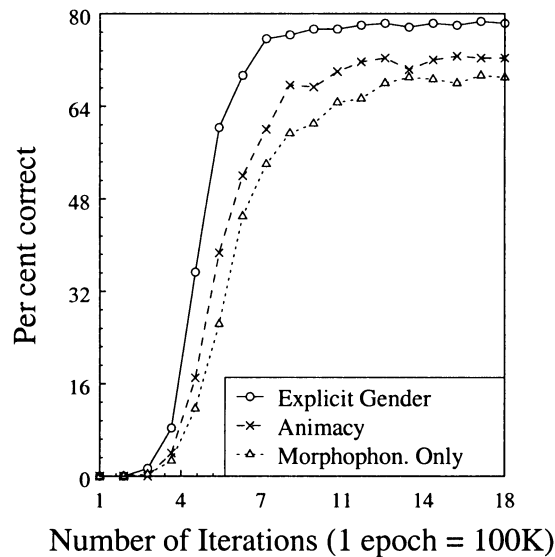


Figure 7. Generalisation test during learning (averaged across 5 runs of each model).

22 MIRKOVIĆ ET AL.

We performed additional analyses on the number of errors in the generalisation test after 1.5 million iterations of learning (approximately 90% of the training corpus correctly produced in all three models) and after 3 million iterations (99–100% of the training corpus correctly produced). There was no significant difference in the pattern of results across the two analyses, so we will present data from the test ran after 1.5 million trials. The analysis of variance on the average number of errors per gender showed significant main effects of model, $F(2, 12) = 5.539, p < .05$, and gender, $F(2, 24) = 18.753, p < .001$; the model \times gender interaction was not significant. Planned comparisons showed that the difference between the Morphophonology Only and Animacy models was not significant. This analysis shows that at asymptote morphophonological information was sufficient to derive novel inflected forms.

In summary, the generalisation test indicates that the models can generalise to new forms, indicating that they had not merely memorised the training sets. The model did not perform perfectly on the generalisation test, but it is difficult to know how this performance would compare to that of human subjects on a similar task (e.g., generating inflected forms for nonce words). The model's performance is likely to undershoot peoples', but the model was exposed to fewer forms and had less capacity to encode semantic regularities.

GENERAL DISCUSSION

The goal of this study was to examine the semantic and phonological correlates of gender in a highly inflected language. Consistent with previous studies of other languages, the corpus analyses indicated that such correlations exist. We identified several, such as the relationships between animacy and abstractness and the masculine and feminine genders shown in Figures 2 and 3. These findings suggest that gender distinctions are strongly correlated with various semantic features. In future research it will be necessary to determine whether there are other semantic correlates of gender, in order to assess the extent to which gender is predictable from multiple probabilistic constraints. This could be achieved by using a connectionist model similar to the one we described above to discover these correlations. For example, a large corpus of words could be assigned featural representations that capture main elements of meaning, as in the McRae et al. (2004) and Harm and Seidenberg (in press) studies. The model could be given these semantic representations as input and trained to classify words as masculine, feminine, or neuter, similar to the use of a connectionist model to discover phonological correlates of adjectives by Haskell et al. (2003). This exercise would provide stronger evidence concerning the extent to which gender is predictable from semantic factors.

The above modelling results are consistent with the corpus analyses insofar as a model lacking an explicit representation of gender both learned to produce correctly inflected forms and generalised to non-trained forms. In future research we anticipate developing models that can extract generalisations about both semantic and phonological correlates of gender and other inflections. We conjecture that the combination of these types of information is highly constraining, but this claim requires further research.

The best performance was achieved in the models that represented gender clearly and unambiguously using dedicated features. It is important to not overinterpret this result. It may be that models that, unlike the ones implemented here, provide a richer representation of lexical semantics would converge on a similar level of performance. The localist model performs better than other models (perhaps too well compared with people) but it cannot encode the correlations between semantics and gender that are present in the language and affect human performance.

Thus, the models provide preliminary data consistent with the hypothesis that grammatical gender can be learned and represented in terms of morphophonological and semantic regularities. This is in accordance with some linguistic analyses (e.g., Corbett, 1991) and with the general PDP approach in which a task such as production is construed in terms of mapping between semantic and phonological representations, with graded representations of gender, case, number and other lexical properties emerging out of learning to master this task.

Challenges for Production Models

Our research is preliminary and none of the results can be taken as definitive evidence against models in which information such as gender is represented by localist nodes at an autonomous level of syntactic representation. As we have noted, further research is needed to determine the extent to which a functional representation of gender arises from phonological and semantic information. That is the main challenge for our approach. However, our research also raises challenges for standard approaches to production. One thing that is prominent in our account and wholly absent from alternatives such as the Levelt et al. (1999) and Caramazza (1997) models is attention to how representations are learned. These representations must be learned because they are language-specific rather than universal. Thus a key question for theories that incorporate abstract syntactic nodes is the mechanism by which a child develops such representations given the information available in the input and facts about the capacities that children bring to the acquisition process. We

know of no account of how a child would create the kinds of abstract, localist syntactic feature nodes that are prominent in standard theories of production. By contrast, in this paper and in other work (e.g., Haskell et al., 2003), we have examined how lexical representations could develop in a system that is learning to perform tasks such as comprehending and producing utterances. From this perspective it is easy to see how a learner would pick up on the semantic and phonological correlates of something like gender; moreover, there is growing evidence that young language learners are also encoding correlational information of this sort (Gerken, 2002). Although we have not demonstrated it here, our strong conjecture is that cross-linguistically, syntactic distinctions such as gender have substantial nonsyntactic correlates and that these are encoded in the course of learning to use language. This information may even prove to be sufficient to support processes such as producing and comprehending correctly inflected forms.

Further progress in deciding between competing accounts of production will require placing them on an even footing. Connectionist models need to address the wealth of behavioural evidence thought to implicate architectures like the ones proposed by Levelt or Caramazza. By the same token, these traditional theories need to consider how the representations incorporated in these architectures could be learned.

Further implications

Although the models presented here were not intended to capture the full range of processes underlying language production (and do not do so), they have implications for current accounts of how gender and other grammatical information is represented. Recall that the two dominant models of word production assume that grammatical information, including gender, number, case, and tense, are represented as abstract nodes that are separate from both semantic and phonological representations (Caramazza, 1997; Levelt et al., 1999). By contrast, we have suggested that these features may be emergent from semantic and phonological information. We presented data from Serbian nouns showing that gender identity is highly correlated with certain semantic properties, consistent with observations in several other languages (Corbett, 1991; Zubin & Köpcke, 1981). We incorporated animacy, the most general of these semantic features, into a model as a first attempt to illustrate how semantic information could guide the computation of a gender-marked phonological form. The clear effects of animacy in the model lead us to speculate that a richer semantic representation may obviate the need for explicit gender and other grammatical nodes in theories of production.

Of course, the viability of an emergent account of gender and other morphological features would not in itself be evidence that abstract grammatical nodes are not a part of the production architecture. That is, explicit nodes might seem rather superfluous, but theories incorporating them could still account for data of the sort that we have presented here. These accounts would have to address other questions to remain viable accounts of language use, however. First, there are now abundant demonstrations of correlations between grammatical features and semantic and phonological representations that affect language users' behaviours (Cassidy et al., 1999; Gerken, 2002; Kelly, 1992) and it is not clear how these correlations could be represented in systems in which gender and other grammatical features are represented through abstract nodes. Second, if gender is represented in terms of abstract syntactic features, how are they learned? Our own view is that gender arises out of semantics and phonology, information the child has to acquire in learning a language. If such information is not sufficient to account for facts about grammatical gender, then how does the child move from knowledge of form and meaning to this abstract level of syntactic representation? Models such as Levelt's and Caramazza's are silent on the issue of where these representations come from.

If gender information is emergent in the mapping between semantic and phonological forms, then it is possible to reconceptualise several important debates concerning the nature of inflectional processing during language production. To this point, we have considered inflectional processing only with regard to the production of individual nouns, but in many languages, features such as number, gender, and case also appear on determiners, adjectives and verbs. In these cases, the inflectional forms depend on the form of a noun; for example in French, a feminine singular noun such as *table* (table) will appear with a feminine singular determiner, any adjective that modifies this noun will also be inflected in the feminine/singular form, and if this noun phrase is the subject of the sentence, the verb will also be marked for singular, as in *La table est petite*. (The table is small.) The computation of agreement during production is typically viewed as a syntactic process in which the abstract grammatical features of the agreement controller (the noun) are passed to other words in the sentence (e.g., Bock et al., 2001). Within this view, a major question concerns interaction between syntax, semantics and phonological representations, specifically the extent to which the syntactic process of feature-passing is influenced by semantic and phonological information. Studies using the fragment completion paradigm have shown that the number of gender agreement errors on predicate adjectives (such as *petite* in the sentence above) is influenced by conceptual gender of nouns (Vigliocco & Franck, 1999, 2001). Vigliocco and Franck found lower error rates when the gender

of the agreement controller had a conceptual basis (as in words for woman, man, nun, father, etc.) compared with the situation in which the noun's gender was more purely grammatical, as in the feminine French noun *table* above. They interpreted these results to indicate that conceptual features can influence the gender agreement process, and similar effects have been shown for number agreement (e.g., Bock, Nicol, & Cooper Cutting, 1999). Number information, like gender, has been assumed to be represented as an abstract syntactic feature. Accounts of agreement production have accommodated influences of conceptual properties of the noun agreement controller (Eberhard, 1999; Vigliocco et al., 1995) but have suggested that the agreement processes should be independent of broader semantic and morphophonological processes (Bock & Miller, 1991; Bock et al., 1999). The most interactive of these models is the Maximalist Input hypothesis (e.g., Vigliocco & Franck, 1999), which suggests that all the available information influences syntactic processes. Importantly, even though these models differ in crucial ways from the non-interactive word production models of Caramazza and Levelt, all of them incorporate the assumption that syntactic information, including grammatical gender and number, is represented on a level separate from semantic and morphophonological information. In this respect, the models differ only in the extent they allow the various levels of representation to interact.

By contrast, if grammatical information like gender emerges from phonological and semantic regularities that exist in the system, then the process of agreement takes on a very different character. The agreement process in this system could be construed in terms of semantic and phonological co-variations between the noun and the agreeing element (determiner, adjective, verb) (see Haskell & MacDonald, 2003; Thornton & MacDonald, 2003, for related discussion). This approach emphasises the important roles of learning and experience in the production process. That is, in the course of learning to comprehend and produce a language, the speaker is sensitive to phonological and semantic co-occurrences. Several previous studies suggest that co-occurrences of this type can modulate production processes. For example, Stallings, MacDonald, and O'Sheghdha (1998) found that the choice of syntactic structure during sentence production was modulated by properties of the main verb in the sentence, which they attributed to speakers' previous experience with verb-structure pairings. Similarly, Dell, Reed, Adams, and Meyer (2000) found that the distributional pattern of phonemes in four-syllable nonsense sequences affected the rate of speakers' speech errors in producing these syllable sequences, and the effects appeared within less than 100 trials.

Applying these general ideas to agreement, it is clear that co-occurrence between elements in a system that is sensitive to sequential patterns should have a powerful effect on production processes. Agreeing elements by

definition co-occur, and a system sensitive to statistical structure will pick up these regularities. For example, in Spanish, the feminine determiner /la/ and nouns ending in /a/ tend to frequently co-occur. Similarly, in semantic terms, agreement controllers (nouns) tend to have a semantic relationship with their agreeing elements; for example subject nouns will tend to be plausible subjects of their agreeing verbs. Thornton and MacDonald (2003) investigated the role of this subject-verb semantic relationship in the production of noun-verb number agreement in English. They found that when another noun in the sentence was also a highly plausible subject, the rate of verb number agreement errors increased. They attributed this result to speakers' sensitivity to distributional patterns in the language, such that noun-verb plausibility relationships become a constraint during the computation of the number marking on the verb. We can speculate that a similar relationship could hold for noun-adjective relationships in languages with gender. On this view, masculine nouns will tend to co-occur with adjectives that denote more male qualities, and feminine nouns will tend to co-occur with more female qualities (e.g., Boroditsky, Schmidt, & Phillips, 2003). This fact can be used in production such that, if semantics is construed in terms of distributed representations, the semantic overlap between the adjective and the noun will facilitate the production of the correct gender-marked phonological forms of both words, because of previous experience of co-occurrence of these phonological patterns.

In summary, the models we have described partially instantiate an alternative perspective on the representation of gender and other grammatical features in language production. This approach emphasizes the acquisition and use of probabilistic information that seems to underlie distinctions such as gender that are treated categorically in most syntactic theories. The research raises several important issues. Are gender nodes necessary or is gender simply a convenient term for a network of probabilistic constraints involving other information, principally semantics and phonology? Are the latter types of information sufficient to account for language processes that involve gender or are other types of representations (e.g., explicit syntactic ones) required? If they are required, and they are not plausibly or demonstrably part of universal grammar, how are they learned? Does the emergentist account of gender extend to other types of grammatical representations, e.g., case, grammatical categories, and other syntactic elements? These questions are likely to be the focus of considerable future research.

REFERENCES

- Bishop, C. (1995). Training with noise is equivalent to tikhonov regularization. *Neural Computation*, 7, 108–116.
- Bock, K., & Eberhard, K. M. (1993). Meaning, sound, and syntax in English number agreement. *Language and Cognitive Processes*, 8, 57–99.
- Bock, K., Eberhard, K. M., Cutting, C. J., Meyer, A. S., & Schriefers, H. (2001). Some attractions of verb agreement. *Cognitive Psychology*, 43, 83–128.
- Bock, K., & Miller, C. A. (1991). Broken agreement. *Cognitive Psychology*, 23, 45–93.
- Bock, K., Nicol, J., & Cooper Cutting, J. (1999). The ties that bind: Creating number agreement in speech. *Journal of Memory and Language*, 40, 330–346.
- Boroditsky, L., Schmidt, L., & Phillips, W. (2003). Sex, syntax, and semantics. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and cognition* (pp. 61–79). Cambridge, MA: MIT Press.
- Caramazza, A. (1997). How many levels of processing are there in lexical access? *Cognitive Neuropsychology*, 14, 177–208.
- Caramazza, A., Miozzo, M., Costa, A., Schiller, N., & Alario, F.-X. (2001). A crosslinguistic investigation of determiner production. In E. Dupoux (Ed.), *Language, brain, and cognitive development: Essays in honor of Jacques Mehler* (pp. 209–226). Cambridge, MA: MIT Press.
- Cassidy, K. W., Kelly, M. H., & Sharoni, L. J. (1999). Inferring gender from name phonology. *Journal of Experimental Psychology: General*, 128, 362–381.
- Coltheart, M. (1981). The MRC Psycholinguistic database. *Quarterly Journal of Experimental Psychology*, 33A, 497–505.
- Corbett, G. G. (1988). Gender in Slavonic from the standpoint of a general typology of gender systems. *The Slavonic and East European Review*, 66, 1–20.
- Corbett, G. G. (1991). *Gender*. Cambridge, UK: Cambridge University Press.
- Cutler, A., McQueen, J., & Robinson, K. (1990). Elizabeth and John: Sound patterns of men's and women's names. *Journal of Linguistics*, 26, 471–482.
- Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26, 1355–1367.
- Eberhard, K. M. (1999). The accessibility of conceptual number to the processes of subject-verb agreement in English. *Journal of Memory and Language*, 41, 560–578.
- Gentner, D. (1981). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology*, 13, 56–83.
- Gerken, L. A. (2002). Early sensitivity to linguistic form. *Annual Review of Language Acquisition*, 2, 1–36.
- Glaser, W. R., & Dungenhoff, F.-J. (1984). The time-course of picture-word interference. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 640–654.
- Harm, M., & Seidenberg, M.S. (in press). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*.
- Haskell, T. R., & MacDonald, M. C. (2003). Conflicting cues and competition in subject-verb agreement. *Journal of Memory and Language*, 48, 760–778.
- Haskell, T. R., MacDonald, M. C., & Seidenberg, M. S. (2003). Language learning and innateness: Some implications of compounds research. *Cognitive Psychology*, 47, 119–163.
- Joanisse, M. F. (2000). *Connectionist Phonology*. (Ph.D. Dissertation. University of Southern California.)

REPRESENTATION OF GENDER 29

- 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.
- Lupker, S. J. (1979). The semantic nature of response competition in the picture-word interference task. *Memory and Cognition*, 7, 485–495.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2004). *Semantic feature production norms for a large set of living and nonliving things*. (Manuscript submitted for publication)
- Miller, G. A., & Fellbaum, C. (1991). Semantic networks of English. *Cognition*, 41, 197–229.
- Mirković, J., Seidenberg, M. S., & Joanisse, M. F. (2004). *Statistical nature of inflectional structure: Insights from a highly inflected language*. (Manuscript in preparation)
- Pearlmutter, B. A. (1995). Gradient calculations for dynamic recurrent neural networks: A survey. *IEEE Transactions on Neural Networks*, 6, 1212–1228.
- Schriefers, H. (1993). Syntactic processes in the production of noun phrases. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19, 841–850.
- Schriefers, H., Meyer, A. S., & Levelt, W. J. M. (1990). Exploring the time course of lexical access in language production: Picture-word interference studies. *Journal of Memory and Language*, 29, 86–102.
- Seidenberg, M. S., & MacDonald, M. C. (1999). A probabilistic constraints approach to language acquisition and processing. *Cognitive Science*, 23, 569–588.
- Stallings, L. M., MacDonald, M. C., & O'Sheghdha, P. G. (1998). Phrasal ordering constraints in sentence production: Phrase length and verb dispositions in heavy-NP shift. *Journal of Memory and Language*, 39, 392–417.
- Stanojčić, Ž, Popović, L., & Micić, S. (1989). *Savremeni srpskohrvatski jezik i kultura izražavanja [Contemporary Serbo-Croatian grammar and style manual]*. Belgrade, Novi Sad: The Institute for Textbooks and Teaching Aids.
- Thornton, R., & MacDonald, M. C. (2003). Plausibility and grammatical agreement. *Journal of Memory and Language*, 48, 740–759.
- Vigliocco, G., Butterworth, B., & Semenza, C. (1995). Constructing subject-verb agreement is speech: The role of semantic and morphological factors. *Journal of Memory and Language*, 34, 186–215.
- Vigliocco, G., & Franck, J. (1999). When sex and syntax go hand in hand: Gender agreement in language production. *Journal of Memory and Language*, 40, 455–478.
- Vigliocco, G., & Franck, J. (2001). When sex affects syntax: Contextual influences in sentence production. *Journal of Memory and Language*, 45, 368–390.
- Vigliocco, G., Hartsuiker, R. J., Jarema, G., & Kolk, H. H. J. (1996). One or more labels on the bottles? Notional concord in Dutch and French. *Language and Cognitive Processes*, 11, 407–442.
- Zubin, D. A., & Köpcke, K. M. (1981). Gender: A less than arbitrary grammatical category. In R. A. Hendrick, C. A. Masek, & M. F. Miller (Eds.), *Papers from the seventeenth regional meeting, Chicago Linguistic Society* (pp. 439–449). Chicago: Chicago Linguistic Society.
- 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* (pp. 139–180). Amsterdam: Benjamins.