

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Semantic Regularities in Grammatical Categories: Learning Grammatical Gender in an Artificial Language

Permalink

<https://escholarship.org/uc/item/54412022>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 33(33)

ISSN

1069-7977

Authors

Mirkovic, Jelena
Forrest, Sarah
Gaskell, Gareth

Publication Date

2011

Peer reviewed

Semantic Regularities in Grammatical Categories: Learning Grammatical Gender in an Artificial Language

Jelena Mirković (j.mirkovic@psych.york.ac.uk)

Sarah Forrest (sf553@york.ac.uk)

M. Gareth Gaskell (g.gaskell@psych.york.ac.uk)

Department of Psychology, University of York
York, YO10 5DD, UK

Abstract

The knowledge of grammatical categories such as nouns and verbs is considered to lie at the foundations of human language comprehension and production abilities. Words' distributional and phonological properties contribute to both adult and infant learning of grammatical categories. Here we investigate the contribution of semantic cues to the acquisition of grammatical categories using grammatical gender. Grammatical gender is traditionally considered a semantically arbitrary category, however there may be finer-grained correlations between semantic categories and gender classes. We taught adult native English speakers an artificial language with two gender-like classes, created via distributional, phonological and semantic properties. We demonstrate that the participants' performance on both an implicit and an explicit task is influenced by the semantic regularities in the two genders. We discuss the implications of the findings for theories of grammatical category learning and use.

Keywords: grammatical categories; grammatical gender; semantic regularities; learning; artificial language

Introduction

How do English speakers know that words like *doctor*, *carrot* and *idea* are nouns and *jump*, *shine* and *think* are verbs? Knowledge of grammatical categories has long been argued to be at the core of our abilities to comprehend and produce language. But how do speakers know that such diverse words have something in common?

Similarities in meanings of words within a grammatical category may be an important cue: in parent-child discourse nouns often refer to objects or persons, while verbs refer to actions (Pinker, 1984). This mapping of meaning to category may be used to "bootstrap" grammatical categories in language acquisition. However, word meaning is probabilistically related to grammatical categories: some nouns are not objects or persons (*idea*), and some verbs are not actions (*think*). Distributional information, such as word co-occurrence, has been suggested as a more reliable source of grammatical category information (e.g. Maratsos & Chalkley, 1980). For example, English nouns are often preceded by *the*, but verbs rarely so. Corpus analyses show that the information about immediately preceding and immediately following words is a reliable cue for categories such as nouns, adjectives and verbs (e.g. Redington, Chater, & Finch, 1998). This and other kinds of distributional information are used by adult and infant language learners as a cue to grammatical category membership (e.g. Gómez, 2002; Mintz, 2002).

Phonological properties also provide valuable information. In English, for example, disyllabic nouns tend to have first-

syllable stress, whereas disyllabic verbs tend to have second-syllable stress; stressed syllables in nouns tend to contain back vowels more than front vowels, unlike verbs (e.g. Kelly, 1992). Grammatical categories in other languages differ along similar dimensions (e.g. Monaghan, Christiansen, & Chater, 2007; Shi, Morgan, & Allopenna, 1998). Processing of words in isolation as well as in a sentential context is influenced by these properties (e.g. Farmer, Christiansen, & Monaghan, 2006).

Importantly, both phonological and distributional properties are probabilistic: disyllabic nouns tend to have first-syllable stress but there are exceptions (*giraffe*, *hotel*). Further, the strength of different cues varies across words. For example, Monaghan, Chater, and Christiansen (2005) demonstrated that in English low-frequency words phonological properties are a more reliable cue to their grammatical category than distributional properties, whereas the opposite is true for high-frequency words.

Here we explore the contribution of probabilistic semantic cues to grammatical category assignment. We focused on grammatical gender, a grammatical category typically considered semantically arbitrary (e.g. Brooks, Braine, Catalano, & Brody, 1993; Caramazza, 1997). For example, *ball* is feminine in Spanish (*la pelota*), but masculine in French (*le ballon*). However, linguistic analyses suggest finer-grained correspondences between semantic categories and grammatical gender (see Corbett, 1991, for a review). For example, in German 70% of nouns denoting alcoholic drinks are masculine, and 77% of nouns denoting reptiles and amphibians are feminine (Corbett, 1991). Moreover, in many languages natural gender is often correlated with grammatical gender (e.g. Spanish: *el hombre_{masc}* - man, *la mujer_{fem}* - woman; German: *der Mann_{masc}*, *die Frau_{fem}*). Again though, there are exceptions (e.g. German: *das Mädchen_{neut}* - girl).

Our aim was to examine the contribution of semantic cues to grammatical gender acquisition. We designed an artificial language with two gender-like classes using distributional and phonological cues, and investigated the contribution of semantic regularities to learning. The effect of semantic regularities on learning what is typically considered a semantically arbitrary grammatical category would provide evidence that semantic information is used even when only probabilistically related to grammatical category membership, and in addition to other available cues.

The artificial language consisted of pronounceable En-

glish pseudo-words, designed to form two gender-like classes based on phonological and distributional properties of the pseudo-words. The items were constructed such that the words in each gender were associated with two different types of endings, and each gender was associated with a specific determiner preceding the words (Table 1).

Table 1: Phonological and Distributional Properties of the Two Gender Classes.

	Determiner	Ending	Example
<i>"Feminine"</i>	tib	eem	tib vedeem
	tib	esh	tib scoiffesh
<i>"Masculine"</i>	ked	ool	ked borchool
	ked	aff	ked thetaff

To create semantic regularities, words in each gender were associated with two different groups of pictures, one representing occupations or characters often associated with different natural genders (e.g. ballerina, sailor), and one representing animals. Occupations and characters depicted as female, and birds, insects and domestic animals were presented with pseudo-words ending in *esh* or *eem* and preceded by the pseudo-determiner *tib* ("feminine"). Occupations and characters depicted as male, and fish, reptiles and wild animals were presented with pseudo-words ending in *aff* or *ool* and preceded by the pseudo-determiner *ked* ("masculine"). The choice of semantic categories was based on semantic regularities in grammatical gender systems in natural languages (Corbett, 1991; Mirković, MacDonald, & Seidenberg, 2005).

We used cross-modal word-picture matching as the training paradigm, combined with pre and post-training cross-modal semantic priming to assess the learning of meanings associated with the pseudo-words (Breitenstein et al., 2007; Dobel et al., 2010). To assess the learning of the gender-like classes, we used two tests. First, we used determiner selection, as a variant of a task often used in language production studies on grammatical gender (Schriefers & Jescheniak, 1999). This test allowed us to assess the learning of gender-like classes based on their association with a distributional cue, similar to the use of gender-marked determiners in languages such as Spanish and French. Second, to test the contribution of semantic regularities to the learning of the two genders, we created an additional set of items (generalization set) with two groups of untrained pseudo-words: *consistent* items conformed to the semantics->phonology mapping of the training items, whereas the *inconsistent* items violated the existing semantics->phonology mapping. Thus in the *consistent* set the items preceded by the determiner *tib* were paired with new pictures of birds, insects, domestic animals and occupations depicted as female, and the pseudo-words preceded by the determiner *ked* were paired with new fish, reptiles, wild animals and occupations depicted as male. In the *incon-*

sistent set, the semantics->phonology mapping was inconsistent with the existing regularities in the training set: items preceded by the determiner *tib* were paired with new fish, reptiles, wild animals and "male" occupations, whereas items preceded with the determiner *ked* were paired with new birds, insects, domestic animals and "female" occupations¹. The generalization set was presented on the final day of training, i.e. when the participants had a good knowledge of the training set. If the participants had learned the semantic regularities associated with the two gender classes, we predicted that their performance should be improved with the generalization items consistent with the previously learned semantics-phonology mapping relative to the inconsistent items. We tested this effect in training using an implicit task (word-picture matching), and post-training using determiner selection as a more explicit task of learning distributional cues associated with different genders.

Experiment

We used a paradigm developed by Breitenstein et al. (2007) to teach participants the new language. This paradigm was originally used to assess the learning of meanings of new words. In both our and the Breitenstein et al. study participants were presented with pictures of objects paired with novel word forms in a probabilistic way: over 5 days each target new word was presented 20 times with a target picture, and twice with a random selection of 10 other (non-target) pictures. The participants' task was to decide whether the word and picture matched, and no feedback was provided. The probabilistic pairing of the new words with the picture targets was intended to reduce any explicit strategies participants may develop during training. In Breitenstein et al., performance improved from 56% correct on Day 1 to 93% on Day 5. To assess how well the new lexical items were integrated with the existing lexical-semantic knowledge, Breitenstein et al. used a cross-modal semantic priming task. Target pictures were preceded by 5 types of primes: new words, picture names, semantically related words, semantically unrelated words and a control set of non-learned pseudo-words. Prior to training, participants were slower to respond to target pictures when they were preceded by the new words relative to semantically related items, and there was no significant difference between the new words and the semantically unrelated items. However, after training the difference between the new words and the semantically related primes disappeared, while the difference between the new words and the semantically unrelated primes emerged. These findings were taken as evidence that the new words had been integrated with the existing lexical-semantic knowledge, and were further confirmed in an MEG study using the N400 effect (Dobel et al., 2010).

We used the same method for teaching the "gender-marked" new words. Before and after learning participants

¹In none of the conditions was the determiner + suffix correspondence violated, i.e. all *tib* items were presented with pseudo-words ending in *esh* or *eem*, and all *ked* items with pseudo-words ending in *aff* or *ool*.

performed semantic priming and determiner selection with the items from the training set. Semantic priming, as well as an English-to-new-word translation recognition task post-training, were used to assess the learning of the meanings of the new words. Determiner selection was used to assess the learning of distributional properties of the gender-like classes. To assess the contribution of semantic regularities, generalization items were presented in addition to the training items on the final day of training, as well as in the determiner selection task subsequent to training.

Method

Participants Eleven undergraduate students from the University of York, UK, participated. All were monolingual native English speakers. They were paid up to £31 depending on their performance on the training task (see below). Two participants failed to complete all training sessions and were removed from the analyses.

Stimuli Training set. The training stimuli consisted of 32 pictures in each gender class, 16 of which depicted occupations or characters, and 16 depicted animals. Birds, insects, domestic animals, and human occupations or characters depicted as female were presented with pseudo-words ending in *esh* or *eem*, which were preceded by the pseudo-determiner *tib*. Fish, reptiles, wild animals and occupations or characters depicted as male were presented with pseudo-words ending in *aff* or *ool*, which were preceded by the determiner *ked*. All pictures were selected from the ClipArt database (Figure 1), and were pre-tested for (English) name agreement with 7 participants who did not participate in the main study.



Figure 1: Example picture stimuli. The picture of the ballerina was paired with *tib scoiffesh*, and the picture of the sailor with *ked borchool*.

The pseudo-determiners (*tib*, *ked*) and the pseudo-suffixes (*esh*, *eem*, *aff*, *ool*) were selected from the ARC database of pronounceable English non-words (Rastle, Harrington, & Coltheart, 2002). Each new word was created by selecting a 3-5 phoneme long English pseudo-word, and attaching one of the suffixes such that the resulting pseudo-word was pronounceable but did not resemble any existing English morpheme or word. Across genders, the words were equated in length, and matched in the beginning and ending phonemes in the stems. Two versions of each item were digitally recorded by a native English speaker (SF), one including the determiner (e.g. *tib scoiffesh*) and the other only including the pseudo-word (e.g. *scoiffesh*). Each version of the item was

recorded twice, such that different recordings were used at training and at test. Items were presented with the determiners during the training, but without them in all other tasks.

Generalization set. Eight additional items were designed for each gender (4 occupations, 4 animals in each gender). Half of the items in each gender were consistent with the training set in terms of the semantics->phonology mapping (e.g. a picture of a female cheerleader was paired with *tib plounesh*, and a picture of a male judge with *ked tarbool*), and half were inconsistent (e.g. a picture of a female teacher was paired with *ked pilmool*, and a picture of a male fisherman with *tib perrbesh*). As in the training set, two versions of each item were recorded, one with and one without the determiner. Two recordings were made for each version, one used on the final day of training, one at test.

Semantic Priming. There were 128 picture targets, of which 64 were used in training (animate targets), and 64 were fillers (inanimate targets). There were four auditory prime types: new words, English picture names, semantically related English words and semantically unrelated English words. For example, for the picture of a ballerina the primes were: *scoiffesh* (new word), *ballerina* (picture name), *dancer* (related) and *crow* (unrelated). All primes were digitally recorded by a native English speaker (SF).

Translation. Sixty-four pseudo-words from the training set and 64 corresponding English words were used in the translation task.

Determiner Selection. The 64 word-picture pairs from the training set were used as the stimuli. We used different audio recordings of the pseudo-words in the pre-training and the post-training tests.

Procedure The schedule of tasks is presented in Table 2. As in Breitenstein et al., to keep participants motivated we provided a monetary reward (additional £1) for each training day that accuracy exceeded 70%. All tasks were implemented in DMDX, and a USB gamepad recorded responses. All items were presented in a fixed random order. Pre-training session lasted 35 min, the training sessions on Days 1-5 20 min, and the post-training session 60 min.

Semantic Priming. Participants were presented with a target picture preceded by an auditory prime and asked to make a binary animacy judgement. Each trial started with a 500 ms centrally located fixation cross, followed by the prime presented over headphones; 300 ms after prime onset a target picture appeared, and it remained on the screen until either the participant responded or a 1500 ms timeout was reached. Experimental stimuli were preceded by 8 practice trials.

Determiner Selection. Participants were presented with a word-picture pair and asked to press a right hand-button if the word was a "tib" word, and a left-hand button if it was a "ked" word (the location of the buttons was counterbalanced across participants). Each trial started with a 500 ms centrally located fixation cross, followed by an auditory presentation of the target word; 200 ms after the target word onset, the associated picture was presented, which remained on the screen

Table 2: Schedule of tasks.

Day	Tasks	Items
Pre-Training (Friday)	semantic priming	T ^a
	determiner selection	T
Training Day 1-4 (Mon - Thurs)	word-picture matching	T
Training Day 5 (Friday)	word-picture matching	T + G ^b
Post-Training (Monday)	semantic priming	T
	determiner selection	T + G
	translation	T

^atraining set; ^b generalization set.

Note: The pseudo-words were presented with the determiners (e.g. *tib scoiffesh*) only during training.

until either the participant responded or a 1500 ms timeout was reached. The two pseudo-determiners were presented on the screen corners corresponding to the response buttons throughout the trial. The experimental stimuli were preceded by 4 practice items.

Word-Picture Matching. Participants were presented with word-picture pairs and asked to press a right-hand button if they matched, and a left-hand button if they didn't. They were encouraged to respond quickly and "based on their intuition, without thinking too much". No feedback was provided. Each trial started with a 500 ms centrally located fixation cross, followed by an auditory presentation a new word; 200 ms post-word onset a picture was presented, which remained on the screen until either the participant responded or a 1500 ms timeout was reached. The experimental trials were preceded by 6 practice items.

As in Breitenstein et al, on each day each new word was presented with the correct picture (e.g. the picture of a ballerina with *tib scoiffesh*) 4 times, and twice with a random selection of two other pictures. Hence over the course of 5 days the correct word-picture pairing was presented 20 times (50%), and twice each with a random selection of 10 other pictures.

Translation. Participants were presented with an English word paired with a new word and asked to press a right-hand button if the two matched, and a left-hand button if they didn't. Each trial started with a centrally located fixation cross which remained on the screen for the duration of the trial; 500 ms post-trial onset an English word was played over the headphones, and 800 ms post-word offset a new word was played. Participants had 1000 ms to respond. Each new word was presented twice: once with the correct English translation, and once with an incorrect English word (one of the words for other correct trials). The experimental trials were preceded by 4 practice items.

Results and Discussion

The time-course of learning. Performance in the word-picture matching task improved over time: the average accuracy on the training set increased from 55% on Day 1 to 83% on Day 5 (Day 2 = 66%, Day 3 = 74%, Day 4 = 80%). An ANOVA with the proportion of correct responses (on the correct word-target picture pairs) as a dependent variable and day (Day 1 or Day 5), gender (masculine or feminine) and semantic category (occupations or animals) yielded a main effect of day ($F(1,8)=17.92$, $p<.01$) and a marginally significant day x gender interaction ($F(1,8)=5.10$, $p=.054$). The interaction indicated that on Day 1 participants were more accurate with masculine targets relative to feminine (66% vs 55%), however this difference disappeared on Day 5 (masculine: 82%, feminine: 85%). On Day 1 the performance was not significantly different from chance in any but the masculine animals condition (one sample t-test; masculine animals: $t(8)=2.91$, $p=.02$; masculine occupations: $t(8)=0.67$, n.s.; feminine animals: $t(8)=-0.54$, n.s.; feminine occupations: $t(8)=1.29$, n.s.). On Day 5 the performance was significantly better than chance in all conditions (feminine animals: $t(8)=5.43$, occupations: $t(8)=3.57$; masculine animals: $t(8)=2.42$, occupations: $t(8)=6.64$; all $ps <.05$).

These findings indicate that the participants learned the items in the training set to a relatively high degree of accuracy.

Learning the meaning of the new words.

Semantic priming. In order to establish whether the participants acquired the meanings of the new words, we tested whether the pre- vs. post-training animacy decisions were influenced by different primes preceding the picture targets. An ANOVA on decision latencies yielded a significant effect of the time of test (pre vs. post-training), $F(1,8)=10.92$, $p<.05$, and prime (new word, picture name in English, semantically related word, semantically unrelated word), $F(3,24)=15.51$, $p<.001$ (Figure 2). Planned contrasts showed that before training, animacy decisions to the pictures were slower after the new words relative to the semantically related English words ($F(1,8)=26.83$, $p<.001$), but this difference disappeared after training ($F(1,8)=0.25$, $p=.63$). This replicates the findings from the earlier studies (Breitenstein et al., 2007; Dobel et al., 2010), and suggests that the new words acquired meanings similar to the semantically related English words.

Translation. We performed an additional test of the acquisition of meaning using the English-to-new-word translation recognition task subsequent to training. The average accuracy rate on this task was over 85%. An ANOVA with the proportion of correct responses as the dependent variable and gender and semantic category as independent variables yielded no significant main effects or interactions (gender: $F(1,8)=0.61$; semantic category: $F(1,8)=0.22$; gender x semantic category: $F(1,8)=0.57$; all ps n.s.). Accuracy rates in all conditions were significantly better than chance (feminine animals: $t(8)=7.57$, occupations: $t(8)=7.33$;




Figure 2: Semantic priming pre- and post-training. (For example, for the target picture *ballerina*: ID = picture name in English; Related = semantically related word (*dancer*), New Word (*scoiffesh*), Unrelated = semantically unrelated word (*crow*).

masculine animals: $t(8)=5.26$, occupations: $t(8)=6.42$; all $ps <.001$). These findings provide additional evidence that the participants had a good knowledge of the meanings of the new words, that was at the same level for both gender classes, and both semantic categories. The performance on this task provides additional evidence of the acquisition of meaning for the new words.

Learning distributional regularities. We tested how well the participants encoded the newly formed grammatical gender classes using a determiner selection task. We compared performance on the same task before and after training in an ANOVA using the proportion of correct responses as a dependent variable and test (pre or post-training), gender and semantic category as independent variables. The accuracy of selecting the correct determiner after training was significantly higher than before training ($F(1,8)=25.18$, $p=.001$), and overall better for occupations than animals ($F(1,8)=9.05$, $p=.017$), Figure 3. Before training, participants' performance in all conditions was at chance (feminine animals: $t(8)=1.07$, occupations: $t(8)=-0.01$; masculine animals: $t(8)=0.42$, occupations: $t(8)=1.72$; all ps n.s.), but significantly better than chance after training (feminine animals: $t(8)=4.90$, occupations: $t(8)=12.60$; masculine animals: $t(8)=6.67$, occupations: $t(8)=7.06$; all $ps <.01$).

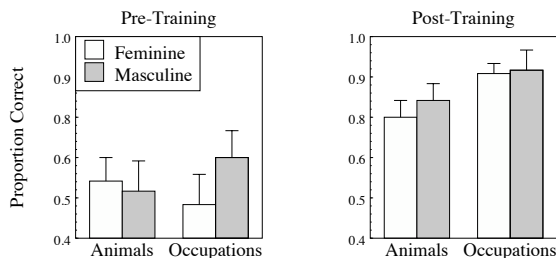


Figure 3: Determiner selection performance on the training items.

These findings suggest that the participants learned a distributional property of the newly formed gender classes, i.e. the co-occurrence between determiners and the 'gender-marked' words. The sensitivity to determiner-noun correspondences in natural languages is often taken as one of the hallmarks of grammatical gender processing in language acquisition and use (e.g. Dahan, Swingley, Tanenhaus, & Magnuson, 2000; Lew-Williams & Fernald, 2007; Van Heugten & Shi, 2009). Our findings suggest that the participants learned the gender-like classes in an artificial language based on their phonological and distributional properties.

Learning semantic regularities. In the final set of tests we explored the contribution of semantic regularities to the learning of the newly formed gender classes. We predicted that new items consistent with the existing semantics->phonology mapping would be easier to process than new items inconsistent with the existing mapping. We tested this hypothesis when the generalization items were first introduced on the final day of training (word-picture matching task), and also after training in the determiner selection task.

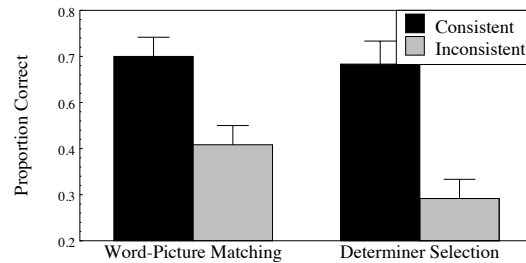


Figure 4: Performance on the generalization items.

As predicted, in the word-picture matching task participants were more accurate with the new items with the semantics->phonology mapping consistent with the training set ($t(8)=2.30$, $p=.05$), Figure 4. Participants' accuracy on the new word-picture pairs consistent with the training items was better than chance ($t(8)=3.83$, $p<.01$), unlike the inconsistent items ($t(8)=-0.56$, n.s.). This suggests that the participants' performance on the new items is facilitated by the existing knowledge of the semantics->phonology correspondences in the training set (cf. Day 1 performance in the training set).

In the post-training determiner selection test, participants were also significantly more accurate in selecting the correct determiner for the consistent relative to the inconsistent items ($t(8)=2.62$, $p<.05$). Interestingly, their performance was significantly better than chance with the consistent items ($t(8)=2.42$, $p<.05$), but significantly worse than chance with the inconsistent items ($t(8)=-2.58$, $p<.05$).

These findings suggest that the participants' performance was influenced by the semantic correlates of the gender classes: the matching of the new item to the correct picture was improved when it was consistent with the existing regularities. The performance on the explicit task assessing the

use of distributional cues further demonstrates the influence of the semantic cues which in this case not only helped when consistent with the existing knowledge, but also hindered performance when the cues were inconsistent with it. It is worth noting that the post-training performance was assessed two days post-training, with no additional training on any items in the meantime.

Conclusions

In this study we taught participants an artificial language with a grammatical gender-like system, using phonological, distributional and semantic cues. Consistent with previous research, participants encoded distributional and phonological regularities of the newly formed gender classes. We demonstrated that their performance is further influenced by the semantic regularities associated with different genders. The use of semantic information was not required by any of the tasks, and was implicit in the stimuli. This finding is particularly relevant given that the category of grammatical gender is traditionally considered semantically arbitrary.

Together the findings provide additional evidence that when learning and using grammatical categories language users rely on probabilistic cues correlated with grammatical category membership, even when they are not in any way useful or necessary for the task at hand.

Acknowledgments

SF was supported by the Undergraduate Summer Bursary from the Department of Psychology, University of York, UK.

References

- Breitenstein, C., Zwitserlood, P., Vries, M. de, Feldhues, C., Knecht, S., & Dobel, C. (2007). Five days versus a lifetime: Intense associative vocabulary training generates lexically integrated words. *Restorative Neurology and Neuroscience*, 25, 493–500.
- 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.
- Corbett, G. G. (1991). *Gender*. Cambridge, UK: Cambridge University Press.
- Dahan, D., Swingle, D., Tanenhaus, M. K., & Magnuson, J. S. (2000). Linguistic gender and spoken-word recognition in French. *Journal of Memory and Language*, 42, 465–480.
- Dobel, C., Junghöfer, M., Breitenstein, C., Klauke, B., Knecht, S., Pantev, C., et al. (2010). New names for known things: On the association of novel word forms with existing semantic information. *Journal of Cognitive Neuroscience*, 22, 1251–1261.
- Farmer, T., Christiansen, M., & Monaghan, P. (2006). Phonological typicality influences on-line sentence comprehension. *Proceedings of the National Academy of Sciences*, 103, 12203.
- Gómez, R. (2002). Variability and detection of invariant structure. *Psychological Science*, 13, 431–436.
- Kelly, M. H. (1992). Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review*, 99, 349–364.
- Lew-Williams, C., & Fernald, A. (2007). Young children learning Spanish make rapid use of grammatical gender in spoken word recognition. *Psychological Science*, 18, 193.
- Maratsos, M. P., & Chalkley, M. A. (1980). The internal language of children's syntax. In K. E. Nelson (Ed.), *Children's language* (p. 127–214). New York: Gardner Press.
- Mintz, T. H. (2002). Category induction from distributed cues in an artificial language. *Memory & Cognition*, 30, 678–686.
- Mirković, J., MacDonald, M. C., & Seidenberg, M. S. (2005). Where does gender come from? Evidence from a complex inflectional system. *Language and Cognitive Processes*, 20, 139–167.
- Monaghan, P., Chater, N., & Christiansen, M. H. (2005). The differential role of phonological and distributional cues in grammatical categorisation. *Cognition*, 96, 143–182.
- Monaghan, P., Christiansen, M., & Chater, N. (2007). The phonological-distributional coherence hypothesis: Cross-linguistic evidence in language acquisition. *Cognitive Psychology*, 55, 259–305.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, Massachusetts: Harvard University Press.
- Rastle, K., Harrington, J., & Coltheart, M. (2002). 358,534 nonwords: The ARC nonword database. *The Quarterly Journal of Experimental Psychology Section A*, 55, 1339–1362.
- Redington, M., Chater, N., & Finch, S. (1998). Distributional information: a powerful cue for acquiring syntactic categories. *Cognitive Science*, 22, 425–469.
- Schriefers, H., & Jescheniak, J. D. (1999). Representation and processing of grammatical gender in language production: A review. *Journal of Psycholinguistic Research*, 28, 575–600.
- 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.
- Van Heugten, M., & Shi, R. (2009). French-learning toddlers use gender information on determiners during word recognition. *Developmental Science*, 12, 419–425.