

Modeling the Effect of Lexico-Syntactic Gender on Spoken-Word Recognition

**Garance M.I. Paris
Matthew W. Crocker**

Department of Computational Linguistics
Saarland University, Germany

Abstract

We present a computational model of the influence of lexico-syntactic gender on spoken-word recognition, and demonstrate the ability of the model to account for relevant findings obtained with eye tracking (Dahan, Swingley, Tanenhaus, & Magnuson, 2000). The model is a Simple Recurrent Network (Elman, 1990) trained on article-noun phrases input phoneme-by-phoneme. It learns to incrementally map this input to object concepts beginning with those sounds. After training, it exhibits a behavior similar to French native speakers using gender to constrain lexical access: When the article preceding a noun is ambiguous in gender, all possible nouns are considered during lexical competition, but when a noun is preceded by a gender-marked article, only nouns belonging to that particular category are considered as potential lexical candidates. In the evaluation, the model is shown to generalize well to novel data including unseen article-noun combinations, leading us to conclude that it has in fact learned an abstract notion of gender and discovered the broader gender patterns in French article-noun sequences.

Keywords: computational model; lexico-syntactic gender; spoken-word recognition; lexical access; eye tracking

Approximately half of the world's known languages subdivide nouns into relatively arbitrary categories known as "*gender*" classes. In these languages, each noun is assigned to a category which is a lexico-syntactic, intrinsic property of the noun itself and often cannot be determined from the noun's form or from its semantics alone. Moreover, depending on the gender of a noun, words that are associated with it change: In French, for example, the singular definite article before masculine nouns is "*le*", but it is "*la*" before feminine ones. Therefore, in principle, after hearing a singular definite article in this language, only half of the nouns in the mental lexicon come into question, because the gender of the noun is foretold by its article. It has been argued, however, that listeners do not make use of gender information in spoken-word recognition because it would be inefficient due to the large number of nouns that would need to be pre-activated (Tanenhaus, Dell, & Carlson, 1987; Jescheniak, 1999). Alternatively, however, pre-activation could be seen as very efficient, since it effectively reduces the search space in the lexicon by half in languages with two gender categories.

Indeed, the bulk of research on gender clearly supports the idea that listeners of gender-marking languages use gender online to facilitate spoken-word recognition. In this paper, we present a model of a mechanism by which gender can constrain lexical access. The model is trained on a corpus of French nouns preceded by singular, gender-marked, and plural (gender-neutral) articles, and learns to simulate the behavior of French natives using gender in spoken-word recognition. Analysis of the model further suggests it does indeed pre-activate nouns based on the article alone. Additionally,

Marshall R. Mayberry, III

School of Social Sciences, Humanities and Arts
University of California, Merced

generalization to unseen article-noun pairs reveals that the network goes beyond learning simple sequential dependencies in the input—as current models of lexical access do—and has learned an abstract notion of gender which influences the earliest stages of lexical access.

Experimental Evidence

Over the past 20 years, findings have consistently demonstrated that speakers of many languages draw on gender information during spoken-word recognition. Both facilitatory and inhibitory effects have been found in several lexical decision experiments in French, Spanish and Russian: Listeners in general were faster at deciding whether a letter or sound sequence was a word or not when the stimuli were preceded by gender-congruent determiners, and slower when they were preceded by gender-incongruent determiners (Grosjean, Dommergues, Cornu, Guillelmon, & Besson, 1994; Faussart, Jakubowicz, & Costes, 1999; Akhutina, Kurgansky, Polinsky, & Bates, 1999). Similar conclusions have also been reached using several other methods (e. g. cross-model priming, Spinelli & Alario, 2002), in other languages (Serbo-Croatian: Gurjanov, Lukatela, Lukatela, Savic, & Turvey, 1985; German: Röder, Demuth, Streb, & Röder, 2003, *inter alia*), with other types of words providing gender information prior to the noun (e. g. demonstratives and possessives, Jakubowicz & Faussart, 1998), in the written modality (i. a. Gurjanov et al.), and recently even with children only 6–7 years old (Roulet, 2007).

Two explanations of the effects have been discussed in the literature by several authors: On the one hand, an interactive-activation model in which the article would pre-activate all congruent nouns, giving them an early advantage over other nouns when they compete for recognition, and on the other hand a post-lexical syntactic congruency-checking mechanism in line with a modular view of lexical access (Grosjean et al., 1994; Bates, Devescovi, Hernandez, & Pizzamiglio, 1996; Friederici & Jacobsen, 1999).

In several recent studies, eye tracking has been shown to be highly sensitive to online lexical access mechanisms. During the recognition of spoken words, listeners are thought to first activate all words matching the onset of the partial input, then gradually eliminate those which become inconsistent with acoustic information. In eye-tracking studies, this appears in participants' eye movements to objects with similar names: For example, Dahan et al. (2000) showed that when gender is of no import, French listeners asked to click on a picture depicting some 'buttons' (*boutons*, /butõ/) also initially looked at 'bottles' (*bouteilles*, /butej) because the ambiguous

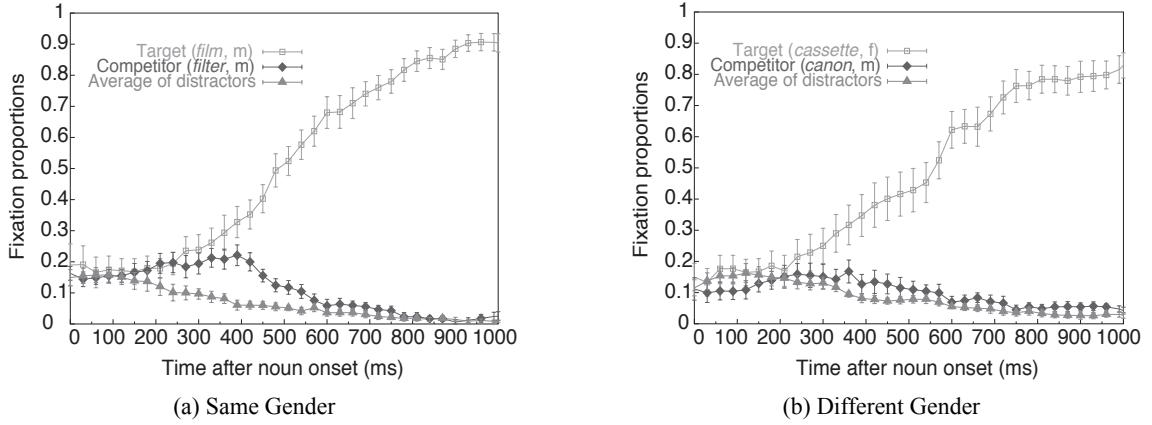


Figure 1: Replication of the Dahan et al. (2000) study using a different design and illustrating two types of eye-tracking curves, with and without competition effect. Statistically, it is the fixations to the competitor and to the averaged distractors which are compared. Competition effects are characterized by more fixations to the competitor than to the distractor, as seen on the left.

noun onset and the neutral French plural definite article allowed both nouns to compete for recognition early on.

However, the picture changed drastically when gender became involved. The authors found that if a singular, gender-marked article came before the noun, the above effect disappeared. Despite the phonological similarity of *bouton* and *bouteille*, when the article preceding the noun in the carrier sentence did not agree with the second picture, the latter was not fixated any more than distractor pictures: When French listeners heard “*le_[masc] bou...[ton]*”, they did not consider the picture of a bottle (*bouteille_[fem]*) because *bouteille* cannot be preceded by “*le*” in French.¹

Importantly, although this finding confirmed results obtained with other methods, it also contributed crucial information about the time-course of the effect and its locus: In the data, it became apparent that gender-inconsistent phonological competitors were eliminated *from the very beginning* of the noun, thereby supporting the idea of an interactive-activation account over a post-lexical congruency check.

Surprisingly however, it also appeared that gender was not drawn upon prior to noun onset. In another group of trials, the authors presented participants with four pictures beginning with different sounds, of which two feminine and two masculine. This time, the statistical analysis revealed no difference between fixations to the distractor belonging to the same gender as the target and to the distractor of the opposing gender. This was taken to mean that gender only begins to influence recognition at noun onset and not earlier, which seems to speak against a pre-activation account in which nouns would be very lightly activated by congruent articles, later giving them an advantage during noun recognition.

Given that the contrast between cases with and without competition will be important for the evaluation of our model, let us now take a closer look at a partial replication of Dahan et

al. (2000). Instead of presenting the same items with the plural and singular definite articles to two groups of participants, two types of trials were displayed to the same participants. In both cases, the target and competitor nouns overlapped at onset, but in the first, target and competitor had the same gender, whereas in the second they did not. The results are presented in Figure 1. When target and competitor were of the same gender, making it impossible for gender to rule out the competitor, a competition effect was observed: At first, all four pictures were fixated equally, then once the noun onset became available, both target and competitor curves started rising while fixations to the distractors decreased.² Finally, when the disambiguating point in the input was reached, the competitor was excluded, and its curve began to drop, rejoining that of the distractor (Figure 1a). By comparison however, when gender information could be used immediately to exclude the competitor, the competitor was not fixated significantly more than the distractors, both curves remaining low all the way throughout (Figure 1b).

Modeling the Effect of Gender

Considering the extreme earliness of the effect in Dahan et al. (2000), it seems that their findings would be most naturally accounted for by an interactive-activation or connectionist model. Well-known models of spoken-word recognition such as TRACE (McClelland & Elman, 1986) and Shortlist (Norris, 1994) can explain a variety of experimental results at the word level and below, but they do not include any mechanism to integrate the influence of context beyond the word. Moreover, their settings are derived in advance by the modelers from linguistic data and pre-wired into the system, without any explanation of how the model might actually acquire such an organization.

¹Similar findings have recently been achieved by Lew-Williams and Fernald (2007) with Spanish-speaking adults and children as young as three years old.

²Due to the time necessary to launch a saccadic eye movement, it is common in eye tracking to observe a 150–200ms delay between acoustic input and fixation data.

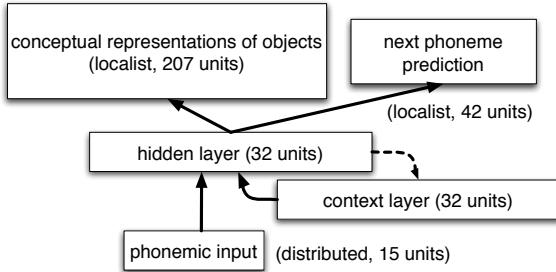


Figure 2: The architecture of the network

In the present research, we were interested in modeling the influence of gender-marked articles on the recognition of subsequent nouns. Therefore, we turned to a type of architecture which is known to be able to deal with syntactic dependencies. Our model is a Simple Recurrent Network (SRN, Elman, 1990, 1991), a type of connectionist network in which the activations of intermediary ‘hidden’ layers are fed back into the network together with the next input. At each step, the hidden units are copied to a context layer and on the next step fed back into the hidden layer, thereby providing the network with a temporal memory and enabling it to learn sequential dependencies in the input.

In addition, SRNs belong to the class of connectionist models which learn from experience. During training, they are presented with examples of input (in our case article-noun sequences), allowed to generate an output, and then corrected by comparing the generated output with the expected one. Based on the discrepancy between the generated output and the correct one, they gradually learn to produce an output which is closer to the desired one.

Network Architecture

The model is trained on article-noun phrases which are input phoneme-by-phoneme, and learns to map the input it has just ‘heard’ with the concepts of objects whose names begin with those sounds. In addition, the network also learns to predict the next phoneme based on the current input, a secondary task which was added following early simulations in which it was found to noticeably improve the performance of the architecture, presumably because it encourages the network to identify sequential dependencies in the input.

On the input group, the representation of the phonemes is distributed, using a feature description adapted from Chomsky and Halle (1968; see the network architecture in Figure 2). The heart of the network consists in a group of 32 hidden units which connectionist networks use to solve their task by recoding the input patterns in an appropriate way.

There are two localist output groups, one for each of the model’s tasks. The first encodes a conceptual representation of the objects contained in the network’s vocabulary, in which each unit stands for one concept. It is here that we can observe what happens during lexical competition: When acoustic input is ambiguous, the model is expected to activate all poten-

tial word candidates corresponding to the input—akin to when participants look at several objects during the ambiguous part of an instruction in an eye-tracking experiment.

On the second output group, the model does next phoneme prediction. Here too, a localist encoding is used. This contrasts with representation of phonemes on the input group, the reason being to facilitate the interpretation of the resulting activations when we evaluate the model (see Evaluation).

Training

The complete corpus used to train the model contained 207 nouns, namely all nouns used by Dahan et al. (2000)³ and an additional 80 nouns added by us. Each noun was paired with the correct definite article (*le* or *la*), the correct indefinite article (*un* or *une*), and the gender-ambiguous plural definite article *les*, giving us a total of 621 patterns, which were split into a training and a test set as follows:

- All nouns were presented to the network with the indefinite article during training.
- In addition, the network saw half of the nouns used by Dahan et al. (2000) in their experiments and all additional nouns (approximately 85 % of all nouns) accompanied by their singular definite article and by the plural definite article.
- The remaining materials from Dahan et al. (2000; 15 % of the total) were held out during training to enable us to later test on unseen singular definite and plural article-noun pairs whether the network learned the generalized concept of gender accurately or not.

We hoped that by presenting 85 % of the nouns with all three articles, the network would have the opportunity to learn the regularity patterns present in the French gender system, i. e. to learn that nouns occurring with *un* can also occur with *le* (and therefore belong to the category which we call “masculine”), that nouns occurring with *une* can also occur with *la* (for feminine nouns), and that all nouns independent of their gender can occur with *les*.

The network was trained using back-propagation through time and bounded momentum descent using the Lens simulator (Rohde, 1999). The complete training materials were presented 200 times during training, the examples from the training file being selected at random by the Lens simulator. The softmax activation function and the cross-entropy error function were used on the output groups to obtain probabilities comparable to the probabilities of fixating each displayed object during an eye-tracking experiment. Successful parameters were identified (learning rate 0.75, momentum 0.9) and the model re-run ten times to ensure its reliability.

³The noun *ballon*, which was used twice in different conditions by Dahan et al. (2000) with different meanings (‘ball’ and ‘balloon’) and therefore depicted by different objects, only occurred once in our materials. Therefore the number of nouns from the Dahan et al.’s second type of trial in Experiment 1 included in our corpus was 63 instead of 64.

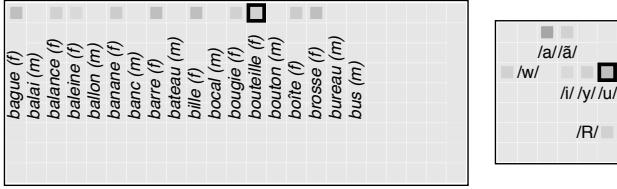


Figure 3: Activations on the concept and phoneme prediction output groups after presentation of the third phoneme in “*la bouteille*” (wrapped vectors of respectively 207 and 42 units). Concepts are sorted alphabetically; therefore nouns beginning with /b/ are grouped in the upper left corner (here labeled with noun and gender). Only feminine /b/-nouns are activated, whereas masculine units are inactive. On the phoneme prediction group, activations closely reflect the probability distribution of phonemes after /b/ in the training corpus.

Evaluation

Figure 3 shows an example of the network performance after training. The model was evaluated by computing the cosine between the network activations and the probability distribution of concepts or upcoming phonemes in the training data, at each phoneme in the input, resulting in a similarity measure for the two vectors contained between 0 and 1.⁴ For the concepts’ group, we expected all possible lexical candidates to be equally activated, assuming a uniform frequency distribution. For example, after presentation of the sequence /lebu/ (*les bou...*), three words in the corpus should remain in the competitor set, gender notwithstanding: *bouton*, *bouteille* and *bougie* (‘candle’). Therefore, each of the three units representing those objects should have an activation of 0.333, and all other units should be inactive. For the phoneme prediction group, we calculated the probability of each phoneme following the current input: After /lebu/ only two possible phonemes remained, /t/ and /ʒ/, but /t/ occurred twice and /ʒ/ only once in the corpus, so that the unit for /t/ was predicted to have an activation of 0.666 and the unit for /ʒ/ of only 0.333.

On the training set, the average cosines over ten runs after 200 training epochs was 0.944 for the concept group and 0.969 for the phoneme prediction group. On the test set (the subset of nouns from Dahan et al., 2000, which were shown to the network only with the indefinite singular or definite plural article during training), the values were 0.887 and 0.975, respectively. This shows that the model generalized well from seen to unseen data: After being trained, it knew that words it had seen only with *un* or *une*, depending on their respective genders, could also occur with *le* or *la*, and that any noun could also be preceded by *les*.

Comparing Model Output and Experimental Data

In addition to cosine evaluation, we also plotted the average activations on the concepts output group when the network was presented with the left-out testing materials from Dahan et al. (2000), averaging over ten runs and using the same counterbalancing of materials as in the original experiments. The

results can be seen in Figure 4. Two kinds of graphs are reported: In the first, the activations from the model were plotted directly, without transformation. The y-axis thus represents the model’s estimation of the probability of each concept being the target mentioned in the input, when taking all 207 nouns in the lexicon into consideration. In the second type of graph, the activations were converted to proportions for the four nouns presented on the screen in the experiment, as is usually done with eye-tracking data.

Overall, the shape of the obtained curves was very similar to what is commonly observed in eye-tracking experiments (see Figure 1): As long as the acoustic input remains ambiguous, all objects are fixated equally often, but as soon as more information becomes available, fixations to the target start rising with every incoming phoneme, while at the same time looks to the distractors decrease. When a competition effect is observed, the competitor rises at first with the target, then falls away and rejoins the distractor curve once the participants can tell which object they are being asked to find.

Figure 4a shows the simulation results for the first type of trial in Dahan et al. (2000), in which the plural, gender-unmarked, article *les* was used. On the first and second phoneme, the article does not allow any distinction between the four pictures, and therefore no difference is present between the activation levels of the four pictures (phoneme 1: mean competitor activation 1.5648×10^5 , mean activation of averaged distractors 2.8612×10^5 , $t(9) = -2.239$, $p > 0.05$; phoneme 2: 2.6268×10^3 , 2.6279×10^3 , $t(9) = -0.032$, $p > .9$). However, the third phoneme is the onset of the noun. At that point, the input is ambiguous between two objects, the target and the competitor, which both start with the same sound, but the distractor is eliminated by the noun onset. Therefore, target and competitor curves start to rise together. This goes on until the end of the overlap, when the competitor begins to drop. The difference between competitor and distractor from phoneme 3–5 is significant (competitor: 0.356086682 , distractors: 1.2952×10^5 , $t(9) = 51.326$, $p < 0.001$). After that, the competitor curve gradually drops due to the averaging over items, because the duration of the overlap in terms of phonemes varies.

By comparison, Figure 4b presents the plots for the trials in which the same nouns were preceded by the gender-marked singular definite article. Again all pictures were activated equally at first (1.5648×10^5 , 2.8612×10^5 , $t(9) = -2.239$, $p > 0.05$), but on the second phoneme, when gender information became available, it is obvious from the proportions in the graph on the right that only the target went up, whereas the activation proportions for the competitor remained similar to those of the distractor: Objects of the opposed gender were thus not activated by the model (mean activations from phoneme 2–5 for the competitor 1.4089×10^2 , for the distractor 1.6407×10^4 , $t(9) = 1.147$, $p > 0.05$).

Finally, let us take a look at what the model did in the case

⁴Although we adopt Elman (1990)’s evaluation method, alternative similarity measures, such as KL-divergence, could also be used.

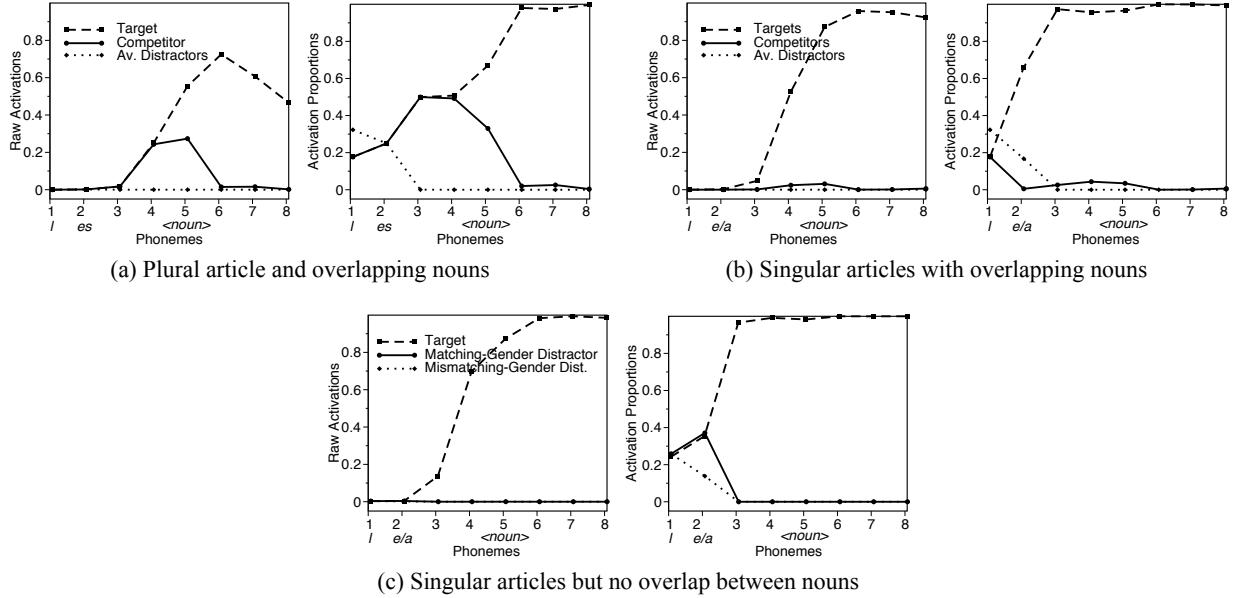


Figure 4: Simulation of the Dahan et al. (2000) experiments: Plots of the average unit activations on the concept output group for the pictures used in the experiments when the model was presented with the same target nouns as input as in the experiments. In each pair of graphs, the left graph presents the raw activations and the right graph the activation proportions, after application of the transformation commonly used in eye-tracking experiments.

of four pictures with non-overlapping names, two of one gender and two of the other (Figure 4c). In the raw activation plot, no difference is visible between gender-matching and gender-mismatching distractors at phoneme two, as was observed by Dahan et al. (2000) in their experiment. However, as is always the case with null-effects, it is also possible that the effect was simply too subtle to be detected statistically. Indeed, when the activations from the model were transformed to proportions, an effect did appear at phoneme 2 which was not previously visible (competitor activation 3.9257×10^3 , distractor 1.4319×10^3 , $t(9) = 4.861$, $p \leq 0.001$).

This is in accordance with the results obtained in a similar eye-tracking study of gender in German conducted by Paris, Weber, and Crocker (2006), in which an adjective was introduced between article and noun, thereby extending the delay between article and noun and making a subtle effect statistically detectable. In this experiment, an activation difference between competitor and distractors was found immediately after article offset, several hundred milliseconds before the onset of the noun. This provides additional support for an account of the gender effect in which gender pre-activates all congruent nouns in the lexicon, giving them a very slight boost, so that they later have an advantage in comparison to other nouns during the process of lexical competition.

It should be noted that the two types of graphs we present appear to be complementary: The similarity with eye-tracking results is more apparent when the activations are plotted straight from the model, but subtle differences are enhanced by the transformation to proportions, thereby highlighting some effects which are not manifest in the plots of raw ac-

tivations. The reason why some differences do not appear in the raw activations (e. g. between competitor and distractor at phoneme 2 in Figure 4b, and between target and the two other pictures at phoneme 2 in Figure 4c) is obvious: Given the number of concepts in the model, the activations are divided by a large number, resulting in near-zero values. This is reminiscent of a frequent argument against eye tracking, which argues that observed effects may crucially be enhanced by the small number of possible visual targets. The comparison between graphs in this study seems to support the refutation offered by proponents of eye tracking, namely that eye tracking does nonetheless reflect more general mechanisms of spoken lexical retrieval.

Discussion

We have presented a model of how gender can constrain lexical access in spoken-word recognition. The model is an SRN, a connectionist architecture which learns from data and has previously been shown to be able to deal with temporal sequences and syntactic relations. The model generalizes well from seen to unseen data after training, allowing the conclusion that it has learned the concept of gender, i. e. knows about the patterns in article-noun sequences in French. In addition, the model closely simulates the results obtained by Dahan et al. (2000).

We would like to emphasize that the model is not intended to compete with specialized models of lexical access, but rather should be seen as the next step toward modeling the online influence of top-down lexico-syntactic constraints on lexical access.

Another interesting characteristic of SRNs and connectionist networks in general which has not been explored here is that they are known to be able to discover classes among the input they deal with and to categorize data. For example, in Elman (1990), it was shown that they could discover part-of-speech categories such as nouns and verbs from the superficial but grammatically structured input they received. In Elman, this was demonstrated after training had been completed by a cluster analysis of the hidden unit activations during a test run. In Elman (1991), the more advanced technique of Principal Components Analysis was used. It would certainly be interesting to apply such techniques to the hidden unit activations of our model to see if they cluster according to gender and if any particular unit or set of units in particular on the hidden layer can be said to encode noun gender, but we leave such an analysis for future work.

Conclusion

In sum, we present a computational model of how lexico-syntactic gender constraints can influence online spoken lexical access, and demonstrate the ability of the model to account for relevant experimental findings. The model simulates the strong effect of gender in lexical competition when onset overlapping competitors are present, but also more subtle effects of pre-activation when there is no overlap, and explains the mixed experimental findings regarding the latter. Importantly, the fact that we observe gender constraints for previously unseen article-noun sequences strongly suggests that the network is generalizing beyond seen phoneme sequences exploited by existing computational models of competitor activation, but rather is learning an abstract notion corresponding to human gender categories that is able to influence the earliest stages of lexical access.

Acknowledgements

The research reported in this paper was supported by IRTG 715 "Language Technology and Cognitive Systems" funded by the German Research Foundation (DFG).

References

Akhutina, T., Kurgansky, A., Polinsky, M., & Bates, E. (1999). Processing grammatical gender in a three-gender system: Experimental evidence from Russian. In Friederici, Garrett, & Jacobsen, 1999.

Bates, E., Devescovi, A., Hernandez, A., & Pizzamiglio, L. (1996). Gender priming in Italian. *Perception & Psychophysics, 85*(7).

Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York: Harper & Row.

Dahan, D., Swingley, D., Tanenhaus, M. K., & Magnuson, J. S. (2000). Linguistic gender and spoken-word recognition in French. *Journal of Memory and Language, 42*.

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

Elman, J. L. (1991). Distributed representations, simple recurrent networks and grammatical structure. *Machine Learning, 7*.

Faussart, C., Jakubowicz, C., & Costes, M. (1999). Gender and number processing in spoken French and Spanish. *Rivista di Linguistica, 11*(1).

Friederici, A. D., Garrett, M. F., & Jacobsen, T. (Eds.). (1999). *Processing grammatical gender* (No. 5). The Netherlands: Springer.

Friederici, A. D., & Jacobsen, T. (1999). Processing grammatical gender during language comprehension. In Friederici et al., 1999.

Grosjean, F., Dommergues, J.-Y., Cornu, E., Guillelmon, D., & Besson, C. (1994). The gender-marking effect in spoken word recognition. *Perception & Psychophysics, 56*(5).

Gurjanov, M., Lukatela, G., Lukatela, K., Savic, M., & Turvey, M. T. (1985). Grammatical priming of inflected nouns by the gender of possessive adjectives. *Journal of Exp. Psychology: Learning, Memory, and Cognition, 11*(4).

Jakubowicz, C., & Faussart, C. (1998, November). Gender agreement in the processing of spoken French. *Journal of Psycholinguistic Research, 27*(6).

Jescheniak, J. D. (1999). Gender priming in picture naming: Modality and baseline effects. In Friederici et al., 1999.

Lew-Williams, C., & Fernald, A. (2007). Young children learning Spanish make rapid use of grammatical gender in spoken word recognition. *Psychol. Sci., 18*(3), 193–198.

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*.

Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition, 52*.

Paris, G. M., Weber, A., & Crocker, M. W. (2006). German morphosyntactic gender and lexical access. In *Proceedings of the 12th annual conference on Architectures and Mechanisms in Language Processing (AMLaP)*. Nijmegen, The Netherlands. (Available from <http://www.coli.uni-saarland.de/~gparis/Publications/>)

Rohde, D. L. T. (1999). *LENS: The light, efficient network simulator* (Tech. Rep. No. CMU-CS-99-164). Pittsburgh, PA: Carnegie Mellon, Dept. of Computer Science.

Roulet, L. (2007). L'accord grammatical de genre dans la dysphasie de développement. *Psychologie française, 52*.

Röder, B., Demuth, L., Streb, J., & Röder, F. (2003). Semantic and morpho-syntactic priming in auditory word recognition in congenitally blind adults. *Language and Cognitive Processes, 18*(1).

Spinelli, E., & Alario, F.-X. (2002). Gender context effects on homophone words. *Lang. and Cognitive Proc., 17*(5).

Tanenhaus, M. K., Dell, G. S., & Carlson, G. (1987). Context effects and lexical processing: A connectionist approach to modularity. In J. L. Garfield (Ed.), *Modularity in knowledge representation and natural-language understanding* (chap. 4). Cambridge, MA: MIT Press.