

Computational Modeling of Assimilated Speech: Cross-Linguistic Evidence

Natalie D. Snoeren (n.snoeren@psych.york.ac.uk)

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

Department of Psychology, University of York

Heslington, York, YO10 5DD UK

Abstract

Models of speech perception have shown divergent views concerning the mechanisms that underlie perception of assimilated speech. In the present study, we evaluated an existing probabilistic model for English place assimilation and conducted a series of simulations on voice assimilated speech in French. The model was trained on a realistic acoustic-phonetic data set of word-final assimilated stops. Our findings showed that the model accommodates the asymmetric assimilation pattern that exists in French, as well as the finding from a cross-modal priming study in French indicating strong regressive context effects for fully voice-assimilated segments. These results suggest that the perceptual system for speech learns to deal with surface variants of speech probabilistically, through experience of phonological variations that are available in the linguistic environment.

Keywords: Speech perception; Cross-linguistic comparisons; Voice assimilation; Place assimilation; Connectionism.

Introduction

A fundamental characteristic of the speech signal is the variability of its phonetic realization. Often, many surface forms map onto a single lexical representation because of the prevalence of such variations. Here we focus on regular alternations described by language-specific phonological rules such as place assimilation in English and voice assimilation in French. For instance, the word-final coronal in *lean* can be produced as something approximating a labial consonant [lim] when followed by a labial context (e.g., *lean bacon*). Because the following context triggers the change in place of articulation, some have argued that listeners rely on the post-assimilation context to derive the underlying identity of the assimilated speech. For instance, Gaskell and Marslen-Wilson (1996) examined the influence of assimilation on lexical access using cross-modal priming. Primes were either assimilated (e.g., *lean*), canonical (e.g., *lean*), or unrelated forms. It was shown that the magnitude of the priming effects was comparable for assimilated and canonical forms in the absence of following context or in the presence of a phonologically viable context (*bacon*). However, in the presence of an unviable context (e.g., *lean gammon*, where the labial place is no longer contextually licensed), priming effects were no longer obtained. The critical role of phonological context led Gaskell and Marslen-Wilson to interpret their results in terms of a language-specific regressive inference mechanism that compensates for assimilation.

There is strong evidence that listeners make use of following context in the resolution of assimilated speech (Gaskell & Marslen-Wilson, 1996, 2001; Coenen,

Zwitserslood & Bölte, 2001; Mitterer & Blomert, 2003). However, the Gaskell and Marslen-Wilson studies only examined complete assimilations that cause a discrete change from one phoneme to another, generated deliberately by a non-naïve speaker. Others, (e.g., Gow, 2002, 2003) focused on assimilation as a graded modification of the speech sound (see e.g., Nolan, 1992, for acoustic-phonetic data). Partially assimilated segments may be considered as ambiguous between two phonemic categories. Gow (2001) performed a phoneme monitoring experiment to determine whether listeners use assimilation to anticipate the segments that trigger assimilation. Gow's study used cross-spliced tokens of spontaneously assimilated speech and found significantly facilitated detection latencies for targets following assimilated segments compared to targets that followed unmodified word forms. Gow (2001, 2002, 2003) proposed that place-assimilated segments always contain information related to both the assimilated and following segment. Hence, listeners should be able to exploit residual cues in the assimilated segment to access its underlying identity. Thus, context information may be perceptually weighted differently according to whether segments are submitted to partial or to full assimilation. Gow (2001) suggested that the reason for the divergent accounts in the temporal processing of assimilated speech might stem from the selection of stimuli representing various degrees of assimilation. This prediction was tested in an empirical study of the perception of voice assimilated speech in French (Snoeren, Seguí, & Hallé, in revision). Cross-modal form priming was used to examine perceptual processing of word-final stop consonants in French that underwent either full or partial assimilations. It was found that for full assimilations following context enhanced perceptual processing, whereas it did not for partial assimilations.

To model the perception of variant speech, Gaskell, Hare, and Marslen-Wilson (1995) proposed a connectionist network that was trained on the mapping between variant speech and invariant, fully specified representations of words. The original model did not incorporate graded assimilation, but given the later demonstrations of the importance of strength of assimilation (e.g., Gow, 2001), the approach was extended to graded assimilation (Gaskell, 2003). The later model correctly reflected the anticipatory and regressive context effects observed in the literature. However, in the absence of realistic data on the prevalence of assimilation of different strengths, the 2003 model was trained using an artificial flat distribution between 0 and 100% assimilation. The model would be more persuasive if it could be shown that similar properties emerge when trained on a more realistic range of assimilation strengths.

The connectionist model described above embodies a “learned compensation” approach. That is, it assumes that the perceptual system learns to deal with surface variants of speech probabilistically, through experience of the particular phonological variations that are available in the linguistic environment. This implies that the perceptual system becomes tuned to the particular language the listener is born into (see, e.g., Mehler, Lambertz, Jusczyk, & Amiel-Tison, 1986; Jusczyk, Friederici, Wessels, Svenkerud, Jusczyk, 1993). It follows that processing of phonological variation is guided by language-specific rules and develops language-specific compensation constraints. This is perhaps the key distinguishing feature of the connectionist model with respect to alternatives that assume a much weaker role for learning. Indeed, Gow and Im (2004) have argued that their feature parsing account is driven by universal perceptual mechanisms, and operates in the same way irrespective of the language environment during development (but see Darcy, Ramus, Christophe, Kinzler, & Dupoux, in press).

However, the learned compensation model (Gaskell, 2003) has only been evaluated with respect to one example of phonological variation from a single language (English place assimilation). Consequently, a key assumption of the model concerning the use of language-specific mechanisms in the processing of regular variations in speech has not yet been addressed. Cross-linguistic comparison is therefore crucial to distinguish between the potential psychological mechanisms that allow listeners to comprehend assimilated speech. In the current simulations we looked at the model’s behavior for French voice-assimilated speech. Given the periodic nature of voicing, voicing cues inherently play out over a longer interval than place cues, which in contrast may often be identified at a single point in time in terms of their spectral acoustic correlates (e.g., Blumstein, 1986). This means that voice cues must be integrated over longer temporal windows than place cues (Gow & Im, 2004). Furthermore, English place assimilation only occurs in one direction, in that coronal consonants can gain velar or labial place but the reverse is not true. French, on the other hand, allows voice assimilation in both directions (although not symmetrically), which has the potential to increase the ambiguity in assimilated speech, and the complexity of the perceptual mapping.

The main objective of the research reported here was to evaluate Gaskell’s (2003) model to examine whether it can accommodate cross-linguistic differences between English place assimilation and French voice assimilation. If a probabilistic model can cope with distinct phonological alternations, this would suggest that the functional difference between the [place] and [voice] features might not be critical during speech perception. We trained the network on compensation for voice assimilation using a realistic French acoustic-phonetic dataset derived from earlier work (Snoeren, 2005; Snoeren et al., 2006).

Properties of voice assimilation in French

Voice assimilation in French can occur when two consonants of different voicing are adjacent (Snoeren et al. 2006). For instance, the final /p/ in *jupe* (“dress”) tends to become voiced when followed by a voiced consonant and *vice versa* (e.g., *une jupe* [jub] *grise*, “grey skirt” vs. *une robe* [rop] *sale*). Snoeren et al. (2006) found that in naturally produced voice assimilated speech, as in *jupe grise* or *robe sale* (“dirty dress”), assimilation is not always complete. More specifically, it was observed that voiceless word-final consonants were perceived mainly as voiced (e.g., /b/ in *jupe*), whereas voiced consonants were perceived less often as voiceless (/b/ than /p/ in *robe*). Acoustic measurements showed that voiced word-final consonants (/b, d, g/) were assimilated to a lesser degree than voiceless word-final consonants (/p, t, k/). Crucially, the correlation between the perceptual judgments and the acoustic measurements was quite high. Because of the graded aspects of French voice assimilation, as well as the observed assimilation asymmetry between voiceless and voiced word-final stops, the current simulation provides a good test case to evaluate Gaskell’s model with respect to the role of the following context and assimilation strength.

Simulating voice assimilation in French

Method

Architecture

The simple recurrent network (SRN; Elman, 1990) used in the present simulations had a similar architecture to the Gaskell (2003) model (see Fig. 1), incorporating a small set of broad features in order to focus on the issue of voicing assimilation strength of French word-final stops. The network was trained using backpropagation to learn the mapping from a surface variation onto the underlying canonical form via recurrent connections at the hidden-unit level, used to accumulate information on temporal sequences. During training, the connection weights between the units in the network were adjusted to facilitate production of the desired mapping. The input consisted of repeated vowel-consonant-consonant sequences, with each segment either being a vowel (V), a consonant that contains voicing (C+) or a consonant that lacks voicing (C-). The possible combinations of these three sequences were the following: VC⁺C⁺, VC⁺C⁻, VCC⁻, VCC⁺ (as in [robgrise], [ropsale], [jupcourte], and [jupgrise])¹. In all cases, presence of a particular segment was marked with a 1 and absence with a 0. At the output level, there were three groups of three units, representing the network’s evaluation of the underlying representation of the previous, current and following segment. Here we focus on the initial evaluation of assimilated segments before the context is known, i.e., as the assimilated segment is presented, based on the current segment output nodes.

¹ The schwa vowel in word-final positions such as in *robe* and in *jupe* is usually deleted in French.




Figure 1: SRN network architecture for the simulations.

We also examined regressive context effects involving a shift in evaluation of a phoneme once its following context is presented to the network, and appearing at the “previous” segment nodes. All simulations were carried out with 20 hidden units intervening between input and output levels, and a further 20 “context” units providing recurrent links.

Training

The training period was intended to allow the model to tune its perceptual evaluation of speech on the basis of exposure to language-specific variation in the phonetic realization of the speech signal. No attempt was made to capture the developmental profile of this tuning—only the end result is important for the current purposes. The task of the network was to provide the output of the canonical form of speech (a representation of the lexical form).

Assimilation distribution

For the purposes of training, tokens were taken from a French acoustic-phonetic dataset containing detailed information about the voicing degree for 160 stimuli (Snoeren et al., 2006). Half of all sequences contained cases in which assimilation occurred (i.e., the presence of two consonants with a different value of the $[\pm\text{voice}]$ feature), the other half contained cases without assimilation (i.e., the presence of two consonants with the same value of the $[\pm\text{voice}]$ feature). The most obvious acoustic cue for voicing in French, as in most other languages, is the presence of vocal fold vibration (cf. Lisker & Abramson, 1964; van Dommelen, 1983), although a number of other cues have been proposed as well (see, e.g., Saerens, Serniclaes, & Beekmans, 1989). Acoustic measurements have shown that the relative duration of the voiced part of the occlusion within the stop consonants provided a consistent index of the presence of voicing for voiceless and voiced words (see Snoeren et al., 2006). Figure 2 shows the distribution of voicing degree according to underlying voicing (12 voiceless words and 8 voiced ones inserted in neutral and post-assimilation contexts). As can be seen from the Figure, there is a striking asymmetry between underlyingly voiceless and voiced stops: voiceless are overall much more affected by assimilation, tending to become voiced. Voiced

stops tend to remain strongly voiced even in the presence of an assimilatory devoicing context.

It is important to note that the perceptual relevance of the acoustic measurements had been assessed in a phonetic categorization task in which participants categorized word-final stops in words extracted from their sentential context (Snoeren, 2005; Snoeren et al., 2006).




Figure 2: Acoustic-phonetic distribution of voicing degree for word-final voiceless stop (white bars) and voiced stops (black bars), that were inserted in neutral or in assimilatory contexts. The distribution shows the frequency of occurrence of a stop in a particular voicing category, based on the acoustic measurements. The distribution is computed over 160 items (40 words x 4 speakers).

Each training word consisted of a three-segment CVC sequence for which the final consonant contained voicing information derived from the acoustic-phonetic database. Degree of assimilation was modeled by allowing variation in the surface or input representation whenever a stop consonant was followed by another stop with different voicing (i.e., C^+C^- or C^-C^+). Full voicing assimilation was represented by $C^- = 1$, $C^+ = 0$; full devoicing assimilation was represented by $C^+ = 0$, $C^- = 1$. All intermediate cases of assimilation strength were implemented by either increasing or reducing the voicing feature. For instance, 60% voicing assimilation would be represented as $C^- = 0.6$, $C^+ = 0.4$, whereas 10% of devoicing assimilation would be represented as $C^- = 0.1$, $C^+ = 0.9$. The network was trained on 6000 sweeps through the training set using the Tlearn simulator (Plunkett & Elman, 1997), with learning rate and momentum at values of 0.1 and 0.3 respectively.

Results

After training, the network performed well on the task of outputting the current segment and the previous segments for unassimilated speech. Here, we focus on the initial response to assimilated speech, and the extent to which assimilation strength modulates regressive context effects.

Initial evaluation of assimilation

The voicing bias in the initial evaluation of assimilated consonants (see Figure 3) was calculated by taking the difference between the activations of the voiced and voiceless output nodes. A value of +1 represents an unambiguously voiced evaluation of the consonant, 0 represents maximal uncertainty as to the underlying identity of the consonant, and -1 represents a clear voiceless evaluation of the consonant.

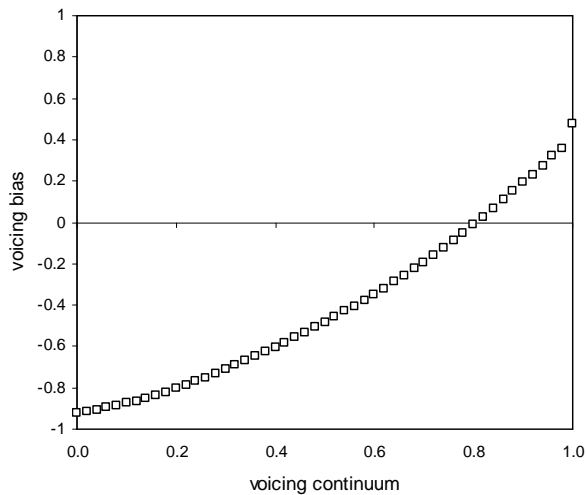


Figure 3: Evaluation of assimilated input before the following segment is identified, represented in terms of a voicing activation bias of the voicing feature as a function of voicing degree. Because following context is not influential at this point in time, each plotted datapoint is the average response of the critical segment followed either by voiceless or voiced contexts.

The curve for the network is fairly straightforward. At low levels of voicing, the model treats consonants as near-unambiguously voiceless ones. This makes sense with respect to the voicing distribution used in training (Fig. 2), since there are far more unassimilated voiceless segments than voiced ones that have undergone strong levels of assimilation in the training distribution. This bias continues towards the middle of the voicing range. However, at high levels of voicing ([80%-100%]), the output of the model remains somewhat ambiguous as to the identity of the underlying consonant, although there is some preference for a voiced consonant. Once again, this behavior reflects the acoustic distribution on which the model was trained, which shows a non-negligible number of voiceless stops (that have been completely assimilated) interspersed with the larger set of unassimilated and weakly assimilated voiced stops.

The model captures the assimilation asymmetry between voiceless and voiced segments quite well. On the one hand, the absence of voicing in the signal unambiguously refers to underlyingly voiceless segments, since voiced segments frequently remain unaffected by assimilation. On the other hand, the presence of voicing gives rise to some ambiguity

as to the identity of the segments since it can both refer to underlyingly voiced segments that have not been submitted to assimilation, or to underlyingly voiceless segments that have turned into voiced ones by voicing assimilation. These results are in accordance with earlier perceptual data that showed that listeners have difficulties in accessing the underlying forms of assimilated voiceless segments whereas they easily accessed those of assimilated voiced segments. This is explained by the fact that the latter stop consonants are less affected by assimilation than the former ones.

Regressive context effects




Figure 4: Effects of following context for critical segments followed by voiceless segments and voiced segments.

The data reviewed in the Introduction showed that the role of post-assimilation context can vary as a function of assimilation strength. Effects of following context in the model were evaluated using the “previous segment” output units, when the context segment was presented to the network (see Figure 4). The graph shows that in the presence of a voiceless context, the model treats a critical segment as voiceless at low levels of voicing and as voiced at high levels of voicing. The evaluation of the critical segment is simply based on the amount of voicing that is present in the segment. This pattern approximates a situation where there is no compensation for assimilation induced by context (as is depicted by the dashed line in Fig. 4). However, the presence of a voiced context increases the tendency for a critical segment to be identified as underlyingly voiceless up to high levels of voicing, leading in this case to compensation for assimilation.

These results, at least when considered at more extreme voicing levels, are in accordance with a recent cross-modal priming data in French (Snoeren et al, in revision). In that study, voiceless stops correspond to complete (or near-

complete) assimilations, whereas voiced ones correspond to partial assimilations. Cross-modal form priming was used, in which assimilated and canonical form primes were auditorily presented to listeners prior to visual lexical decision of the target word. The visual target always corresponded to the canonical word. As can be seen from Figure 5, the priming effect of assimilated voiceless-stop words (e.g., *jupe*) was significantly enhanced when post-assimilation context was made available to listeners, whereas it was not for voiced-stop words (e.g., *robe*).

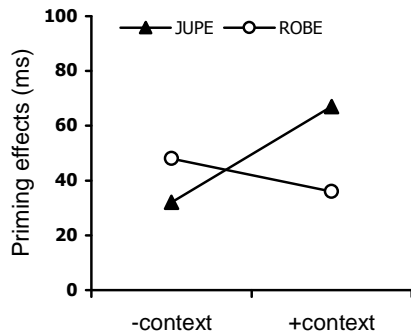


Figure 5: Priming effects for targets words with a voiced vs. voiceless offset (e.g., *robe* vs. *jupe*) primed by assimilated word forms. The absence or presence of context in the auditory prime is noted “-context” vs. “+context” (taken from Snoeren et al., in revision).

These priming results suggest that listeners rely heavily on context for underlyingly voiceless segments, submitted to (near-) complete assimilation, whereas such was not the case for voiced stops that were submitted to low or moderate levels of assimilations.

The present simulation work shows that, according to the probabilistic model that increasingly develops sensitivity to language specific phonological variation; the presence of a voiceless context does not really change the evaluation of underlyingly voiced segment. Inversely, the presence of a voiced context does enhance the recovery of the voiceless identity of the critical segment at high levels of voicing. This pattern of results is in line with the current existing data on the perceptual processing of French voice assimilated speech.

General Discussion

In the present research, we have evaluated Gaskell’s (2003) probabilistic model that embodies the assumption that the perceptual system learns to compensate for assimilatory alternations as a consequence of exposure to regular variations in speech. The model was originally designed to account for a range of temporal effects on the perception of place assimilated speech in English. Here, the model was evaluated in terms of a cross-linguistic comparison while focusing on the model’s behavior regarding voice assimilated speech in French.

We trained the model on a realistic acoustic-phonetic data set of French voice assimilated speech (taken from Snoeren, 2005, and Snoeren et al., 2006). In spite of the functional differences between place assimilation in English and voice assimilation in French, our modeling results are promising. In its initial evaluation of assimilation, the model was able to accommodate the assimilation asymmetry that exists between voiceless and voiced stop consonants in French. At low levels of voicing, there was a tendency to treat the critical segment as unambiguously voiceless, whereas at high levels of voicing, the model showed uncertainty with respect to the identity of the critical segment. This pattern was closely paralleled by the acoustic data set of voicing with which the model was trained. Moreover, a context effect was observed at high levels of voicing for critical segments that are followed by a voiced context segment, which is in agreement with recent data from a priming study in French (Snoeren et al., in revision).

The current simulations fare well with the existing production and perception data in French and are in line with the view according to which learnt language-specific mechanisms underlie listeners’ ability to cope with assimilation and other types of regular variation in speech.

One other aspect of the learned compensation model pertains to the predictive value of assimilation for moderate or mild assimilations. In the present simulation work, we focused on initial evaluation of assimilation and regressive context effects, because of the available empirical evidence in French so far. Future behavioral work should look into anticipatory processing of partially assimilated segments in French to test the profile of predictive strength of assimilation.

Complementary approaches to modeling assimilation

In the present paper, we have examined probabilistic connectionist modeling of regular variations in French and English. This type of modeling encapsulates the assumption that the informational value in strength of assimilation is learnt from experience in a particular linguistic environment, and interpreted probabilistically, rather than dichotomously. In fact, the connectionist implementation could be just as easily replaced by a purer Bayesian model of assimilation perception. Vision is probably the subfield of cognitive science where probabilistic models are most advanced. However, the Bayesian approach has already proved to be very useful to explain various aspects of language processing (see, e.g., Chater & Manning, 2006), suggesting that probabilistic methods have a range of valuable roles to play in understanding human cognition.

Initial work using the Bayesian approach (estimating the likelihood with which a stop consonant was intended as a voiced or as voiceless, given the prior probabilities of occurrence of this particular feature in the French language) showed much the same pattern of results. Indeed, the retention of a connectionist framework in this case is simply to increase comparability between current and previous simulation results. Irrespective of implementation, the

crucial components are the accurate characterization of assimilation in production, and the availability of clear behavioral data in order to assess the performance of the model. In the case of French voice assimilation, we now have all of these components in place, and the broad fit between data and model is good. Such findings strengthen our proposal that the ability to cope with regular variation in speech is likely to be driven by language-specific learning. The resultant probabilistic mechanisms cause particular sensitivity to voice assimilated speech for French listeners and to place assimilated speech for English listeners.

Acknowledgments

This work was supported by a research grant from the British Royal Society and a post-doctoral research fellowship from the French Fyssen Foundation.

References

- Blumstein, S.E. (1986). On acoustic invariance in speech. In J.S. Perkell & D.H. Klatt (Eds.), *Invariance and variability in speech processes*, (pp. 178-201), Hillsdale, NJ: Erlbaum.
- Chater, N. & Manning, C.D. (2006). Probabilistic models of language processing and acquisition. *Trends in Cognitive Science*, 10, 335-344.
- Coenen, E., Zwitserlood, P. & Bölte, J. (2001). Variation and assimilation in German: Consequences for lexical access and representation. *Language and Cognitive Processes*, 16, 535-564.
- Darcy, I., Ramus, F., Christophe, A., Kinzler, K., & Dupoux, E. (In press). Phonological knowledge in compensation for native and non-native assimilation. In: F. Kügler, C. Féry & R. van de Vijver (eds.) *Variation and Gradience in Phonetics and Phonology*. Mouton De Gruyter.
- Dommelen, W. A. van (1983). Parameter interaction in the perception of French plosives. *Phonetica*, 40, 32-62.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Gaskell, M.G., Hare, M., & Marslen-Wilson, W.D. (1995). A connectionist model of phonological representation in speech perception. *Cognitive Science*, 19, 407-439.
- Gaskell, M.G., & Marslen-Wilson, W.D. (1996). Phonological variation and inference in lexical access. *Journal of Experimental Psychology: Human Perception & Performance*, 22, 144-158.
- Gaskell, M.G. & Marslen-Wilson, W.D. (1998). Mechanisms of phonological inference in speech perception. *Journal of Experimental Psychology: Human Perception & Performance*, 24, 380-396.
- Gaskell, M. G. & Marslen-Wilson, W. D. (2001) Lexical ambiguity resolution and spoken word recognition: bridging the gap. *Journal of Memory and Language*, 44, 325-349.
- Gaskell, M.G. (2003). Modelling regressive and progressive effects of assimilation in speech perception. *Journal of Phonetics*, 31, 447-463.
- Gow, D.W. (2001). Assimilation and anticipation in continuous spoken word recognition. *Journal of Memory and Language*, 45, 133-159.
- Gow, D. W. (2002) Does English coronal place assimilation create lexical ambiguity? *Journal of Experimental Psychology: Human Perception and Performance*, 28, 163-179.
- Gow, D.W. (2003). Feature parsing: Feature cue mapping in spoken word recognition. *Perception & Psychophysics*, 65(4), 575-590.
- Gow, D.W. (2003). How representations help define computational problems: Commentary on Grossberg, Gaskell and Greenberg. *Journal of Phonetics*, 31, 487-493.
- Gow, D.W. & Im, A.M. (2004). A cross-linguistic examination of assimilation context effects. *Journal of Memory and Language*, 51, 279-296.
- Jusczyk, P.W., Friederici, A.D., Wessels, J.M.I., Svenkerud, V., & Jusczyk, A.M. (1993). Infants' sensitivity to the sound pattern of native-language words. *Journal of Memory and Language*, 32, 402-420.
- Lisker, L., & Abramson, A.S. (1964). A cross-language study of voicing in initial stops: Acoustic measurements. *Word*, 20, 324-422.
- Mehler, J., Lambertz, G., Jusczyk, P. & Amiel-Tison, C. (1986). Discrimination de la langue maternelle par le nouveau-né. *C.R. Académie des Sciences*, 303, Série III, 637-640.
- Mitterer, H. & Blomert, L. (2003). Coping with phonological assimilation in speech perception: Evidence for early compensation. *Perception & Psychophysics*, 65, 956-969.
- Nolan, F. (1992). The descriptive role of segments: Evidence from assimilation. In G.J. Docherty & D.R. Ladd (Eds.), *Laboratory Phonology II: Gesture, segment, prosody*, pp. 261-280). Cambridge: Cambridge University Press.
- Plunkett, K. & Elman, J.L. (1997). *Exercises in rethinking innateness: A handbook for connectionist simulations*. Cambridge, MA: MIT Press.
- Saerens, M., Serniclaes, W., & Beekmans, R. (1989). Acoustic versus contextual factors in stop voicing perception in spontaneous French. *Language and Speech*, 32, 291-314.
- Snoeren, N.D. (2005). Variations phonologiques en production et perception de la parole: Le phénomène de l'assimilation. Doctoral dissertation, Department of Psychology, University Paris – Descartes (Paris 5).
- Snoeren, N.D., Hallé, P.A. & Seguí, J. (2006). A voice for the voiceless: Production and perception of assimilated speech in French. *Journal of Phonetics*, 34, 241-268.
- Snoeren, N.D., Seguí, J., & Hallé, P.A. (In revision). Perceptual processing of fully and partially assimilated words in French. *Journal of Experimental Psychology: Human Perception & Performance*.