Recognition of signed and spoken language: Different sensory inputs, the same segmentation procedure

Eleni Orfanidou [a,*], Robert Adam [a], Gary Morgan [a], James M. McQueen [b]

[a] Department of Language Communication Science, City University, London, UK
[b] Max Planck Institute for Psycholinguistics and Radboud University Nijmegen, Nijmegen, The Netherlands

ARTICLE INFO

ABSTRACT

Signed languages are articulated through simultaneous upper-body movements and are seen; spoken languages are articulated through sequential vocal-tract movements and are heard. But word recognition in both language modalities entails segmentation of a continuous input into discrete lexical units. According to the Possible Word Constraint (PWC), listeners segment speech so as to avoid impossible words in the input. We argue here that the PWC is a modality-general principle. Deaf signers of British Sign Language (BSL) spotted real BSL signs embedded in nonsense-sign contexts more easily when the nonsense signs were possible BSL signs than when they were not. A control experiment showed that there were no articulatory differences between the different contexts. A second control experiment on segmentation in spoken Dutch strengthened the claim that the main BSL result likely reflects the operation of a lexical-viability constraint. It appears that signed and spoken languages, in spite of radical input differences, are segmented so as to leave no residues of the input that cannot be words.

Crown Copyright © 2009 Published by Elsevier Inc. All rights reserved.

Introduction

Lexical segmentation is an essential step in language comprehension: The perceiver must recognize the words in the incoming message, and to do so must segment the input signal into its discrete lexical components. In reading an alphabetic script such as this, this segmentation problem is easily solved: The reader can use the white spaces as markers to word-boundary locations. But both spoken and signed languages are continuously coarticulated, with no reliable word-boundary markers that are equivalent to the white spaces that appear consistently in printed text. Unlike input using artefactual linguistic transcriptions, therefore, natural-language input brings with it a nontrivial segmentation problem. We ask here whether this problem is solved in the same way in speech and sign comprehension.

On the one hand, one might expect there to be fundamental differences in how spoken and signed language are recognized. Signed languages are produced through the simultaneous use of various upper-body articulators (e.g., hands, arms, fingers) and are perceived in the visual modality. This means that, although signs unfold continuously over time (just like speech), the sign comprehender receives information in parallel about different sign parameters (e.g., handshape, location and movement; Brentari, 1998; Stokoe, 1960). Speech is produced by coordinated movements of vocal tract articulators and perceived primarily in the auditory modality (though there are also visual cues such as lip movements). Although speech information arrives in parallel (e.g., formant transitions in the early part of a vowel can signal both what that vowel is and what the preceding consonant was, Stevens, 1998, and information about e.g. voicing and place of

* Corresponding author. Address: Department of Cognitive, Perceptual and Brain Sciences, Deafness, Cognition and Language Research Centre, University College London, 49 Gordon Square, London, WC1H 0PD, UK. Fax: +44 (0)20 7679 8691.

E-mail addresses: e.orfanidou@ucl.ac.uk (E. Orfanidou), Robert.Adam.1@city.ac.uk (R. Adam), G.Morgan@city.ac.uk (G. Morgan), James.McQueen@mpi.nl (J.M. McQueen).

articulation can arrive simultaneously), spoken words unfold over time in a more serial fashion than signs do. That is, while a spoken word can be characterized as a sequence of partially overlapping consonants and vowels (this primarily sequential information source is labeled "horizontal" by Brentari (1998)), a sign can be characterized as a more simultaneous bundle of parameters (as "vertical"; Brentari, 1998). As a consequence, lexical transitions also differ. A transition from one speech sound to another could be a within- or a between-word transition, but in sign language transitions within words can be qualitatively different from those between words (e.g., the movement within a sign located on the torso can be very different from the transitional movement from that sign to one on the face). These transitional movements, though less reliable and much more variable than the white spaces in text, could nevertheless signal word boundaries in sign. Preliminary evidence suggests that BSL signers do indeed use sign transitions in segmentation (Orfanidou, Adam, McQueen, & Morgan, 2008). Strings of primarily "sequential" spoken words could therefore be segmented very differently from strings of primarily "simultaneous" signs.

On the other hand, however, sign and speech input pose the same computational problem. That is, in the Marrian sense (Marr, 1982), the same computation must be carried out in each case. In spite of the radical modality differences, the perceiver (the sign comprehender or listener) must segment a discrete sequence of lexical units out of a quasi-continuous input which unfolds over time. It is possible, therefore, that the same segmentation procedures are applied across modalities. If this could be shown, then it would be possible to posit language-general segmentation processes rather than those which are limited to speech.

We tested here whether segmentation processes in speech and sign are indeed the same. Much more is known about speech recognition than about sign recognition. Our focus was thus on a theory derived from research on spoken language. It has been proposed that listeners segment continuous speech so as to leave no residues that themselves are not viable words (Norris, McQueen, Cutler, & Butterfield, 1997). For example, in the utterance "Is sign like speech?", segmentations including sigh n would be dispreferred, because nonsyllabic n is not a possible English word. More specifically, according to the Possible Word Constraint (PWC), a lexical hypothesis is disfavored if the stretch of speech between that hypothesis and a likely word boundary does not contain a vowel. The PWC is an essential algorithm in a broader theory of word segmentation and recognition (Norris, 1994; Norris & McQueen, 2008). This theory has two other primary components. The first is that word recognition is based on the competitive evaluation of multiple lexical hypotheses (Luce & Pisoni, 1998; McClelland & Elman, 1986; McQueen, Norris, & Cutler, 1994). The second component derives from the fact that, while speech has no singular reliable segmentation cue like the white spaces in printed text, it nevertheless does contain many cues to the location of likely word boundaries. Silence, and phonotactic, rhythmic and fine-grained phonetic cues can all indicate where word boundaries might be, and listeners appear to use such cues in segmentation (Cho, McQueen, & Cox, 2007; Cutler & Norris, 1988; Davis, Marslen-Wilson, & Gaskell, 2002; Mattys, White, & Melhorn, 2005; McQueen, 1998; Norris et al., 1997). The PWC was proposed as an algorithm which would bring together these two components (lexical competition and cue use) in a unified account of continuous-speech segmentation. Via the PWC, cues to likely word boundaries modulate the competition process so as to rule out spurious lexical hypotheses (such as sigh in "Is sign like speech?", with the impossible word "n" between the offset of sigh and the likely word boundary between the [n] and the [l], cued e.g. by the phonotactic restriction that [nl] cannot occur within a syllable; McQueen, 1998).

The PWC helps to solve the segmentation problem by disfavoring lexical parses which include impossible words. It also helps the listener segment speech that includes new words (Norris & McQueen, 2008; Norris et al., 1997). Armed with the PWC, the sign comprehender could reap the same benefits: Parses with impossible signs could be avoided, and new signs could be recognized more easily. These functional benefits make the PWC a good candidate for a segmentation algorithm that would apply across modalities. But is there a lexical-viability constraint in sign-language segmentation?

We attempted to answer this question by studying, for the first time, segmentation in British Sign Language (BSL). The study of the phonological structure of soundless languages such as BSL has revealed that they have a level of linguistic organization based on minimal meaningless parameters. The location parameter specifies where the hand is located in space in relation to the body (e.g., chest, nose, head, chin), the movement parameter describes how the hand moves in the sign space (e.g., arc, circle, wiggling fingers), and the handshape parameter indicates the form the hand itself takes (e.g., fist, flat hand or index-finger pointing). These parameters combine in rule-governed ways to create lexical signs with meanings (Brentari, 1998; Stokoe, 1960; Sutton-Spence & Woll, 1999). For example, the BSL sign meaning NAME uses the same handshape and movement parameters as the sign AFTERNOON, but they differ in location. Sign parameters are also implicated in sign comprehension and production. In particular, different parameters appear to contribute in different ways to sign processing (Carreiras, Gutierrez-Sigut, Baquero, & Corina, 2008; Corina & Emmorey, 1993; Corina & Knapp, 2006; Dye & Shih, 2006; Emmorey & Corina, 1990; Orfanidou, Adam, McQueen, & Morgan, 2009). This evidence about parameters, together with evidence about larger linguistic units in sign processing, indicates that signed languages have multiple levels of linguistic organization (Sandler & Lillo-Martin, 2006). This suggests in turn that signed languages are fully-fledged natural languages.

Sign languages also have systems of phonotactic constraints which determine possible combinations of handshapes, movements and location changes (Brentari, 1998). For example, all well-formed signs must have a movement, just as spoken words have syllabic nuclei. For signs in American Sign Language with two movements, a number of constraints on the combinations of movement parameters have been proposed (Brentari, 2006; Uyechi, 1996). For example, combinations of straight and arc

movements, and of path movements with a sign-internal movement (i.e., a handshape or orientation change) are not permissible. Similar phonotactic constraints appear to hold in BSL (Orfanidou et al., 2009). These constraints were used in designing the materials in Experiment 1, and especially the impossible signs.

The experimental task was a sign-language variant of the word-spotting task (Cutler & Norris, 1988; McQueen, 1996). In word spotting, listeners hear a list of nonsense sequences, press a response button when they detect a real word embedded in the nonsense, and then say aloud the word they spotted. This task has provided the primary evidence for the PWC in speech across a wide variety of spoken languages. For example, English listeners spotted words such as *apple* more slowly and less accurately in *fapple* than in *vuffapple* (Norris et al., 1997). In the former sequence there is a nonsyllabic residue between the initial silence and the target's onset (the single consonant *f*, an impossible English word); the PWC theory correctly predicted that spotting *apple* in *fapple* would be difficult. But in the latter sequence the equivalent residue is a syllable (*vuff* is not an English word, but could be), and, following the PWC predictions, word spotting was easier. The word-spotting task and this manipulation of impossible- and possible-word contexts were adapted here to BSL. Deaf BSL comprehenders saw nonsense sequences consisting of two signs. On some trials, the second sign was a real BSL sign. The participants' task was to press a button when they spotted BSL signs, and then to sign them to a video camera. Targets appeared after a nonsense sign that was either a possible or an impossible BSL sign. If viewers of sign, like listeners to speech, are sensitive to whether component parts of the incoming linguistic signal are possible signs, then they should find it easier to spot signs in possible- than in impossible-sign contexts.

We also manipulated Age of Acquisition (AoA). A listener's first spoken language is usually acquired from birth. The situation in the Deaf community is very different. Only 5–10% of deaf children are born into Deaf families, and early acquisition of a sign language is rare. The majority of Deaf people's first exposure to other signing peers is at school, and for some it is even early adulthood. The effect of late exposure to sign language on comprehension abilities appears to be negative and long lasting (Corina & Hildebrandt, 2002; Mayberry, 1994; Mayberry & Eichen, 1991; Newport, 1990). AoA does not appear to influence all aspects of sign processing, however (Newport, 1990). For example, offline tests of basic word order and a grammaticality judgment task show no effects of AoA (Emmorey, Bellugi, Friederici, & Horn, 1995; Mathur & Best, 2007; Newport, 1990). We compared signers who acquired sign in the first few years of life with those acquiring sign in childhood and in adolescence. On the one hand, effects of AoA on phonological processing (Corina & Hildebrandt, 2002; Mayberry & Eichen, 1991) might suggest that signers with earlier AoA would be more likely to exhibit use of the PWC algorithm in their segmentation behavior. On the other hand, if the PWC is a valuable segmentation algorithm for all sign comprehenders, there may be no AoA effect.

There is, however, a fundamental difference in what can constitute an impossible word between signed and spoken languages. In speech, impossible words are vowelless sequences such as single consonants. Because spoken words are made by stringing sequences of vowels and consonants together, impossible words arise when essential phonological material (i.e., a vowel) is missing. This can never occur in BSL. Because of the physical nature of sign articulations, and the simultaneous nature of sign-language phonology, every BSL sign must have a handshape, a location, and movement (minimally, movements from and to adjacent locations). It is thus not possible to remove parameters to create impossible signs. We created them instead by adding superfluous phonological parameters to existing signs, resulting in nonexistent, impossible combinations. In Experiment 1, impossible-sign contexts with too many parameters were compared with possible-sign contexts with legal (but nonexistent) parameter combinations.

Experiment 1

Method

Participants

Sixty Deaf BSL signers between the ages of 18 and 60 years took part. There were 34 native Deaf BSL signers (exposed to sign before 5 years of age), 13 childhood BSL learners (exposed to sign between 6 and 12 years of age), and 13 adolescent BSL learners (exposed to sign after 12 years of age). All had normal or corrected visual acuity and were paid to take part.

Stimuli

The targets were the 32 simple BSL signs listed in Appendix A. They were each embedded in two types of preceding nonsense context: a possible nonsense sign (i.e., a sign that does not exist in BSL but could exist) and an impossible nonsense sign (i.e., a sign that could not occur in BSL). A Deaf linguist who is also a native signer (the second author) created two hundred nonsense items made up of possible and impossible signs which were then coded and filmed. The nonsense signs were made by combining phonological parameters (i.e., handshape, movement, location and orientation) in novel ways. We included marked and unmarked handshapes (Sutton-Spence & Woll, 1999), major body areas (head, trunk, neutral space, nondominant hand and specific locations within these major areas) and major movements (straight, arc, circle). An example of a nonsense possible sign followed by a real BSL sign (LAZY) is shown in Fig. 1. The nonsense sign has a straight movement up the forearm ending in a closed handshape at the nondominant hand. There is no change in orientation.

The impossible nonsense signs were illegal combinations of phonological parameters, for example, straight and arc movements combined with a hand-internal movement. An example of a nonsense impossible sign followed by a real BSL sign (MOON) is shown in Fig. 2. The impossible sign moves forward then upwards in an arc movement. As the hand makes the arc it changes orientation (palm facing upwards to downwards) and the handshape changes to an open hand.
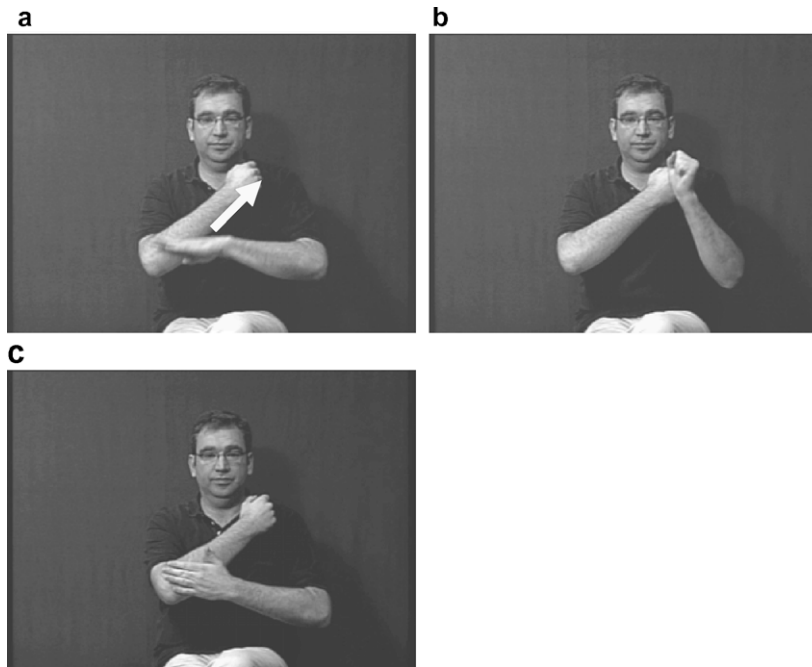
Fig. 1. A possible nonsense sign and target in three video stills. In the possible nonsense sign the hand first moves up the forearm (a). It closes at the nondominant hand (b), before the transition to the target (the BSL sign LAZY; (c)).
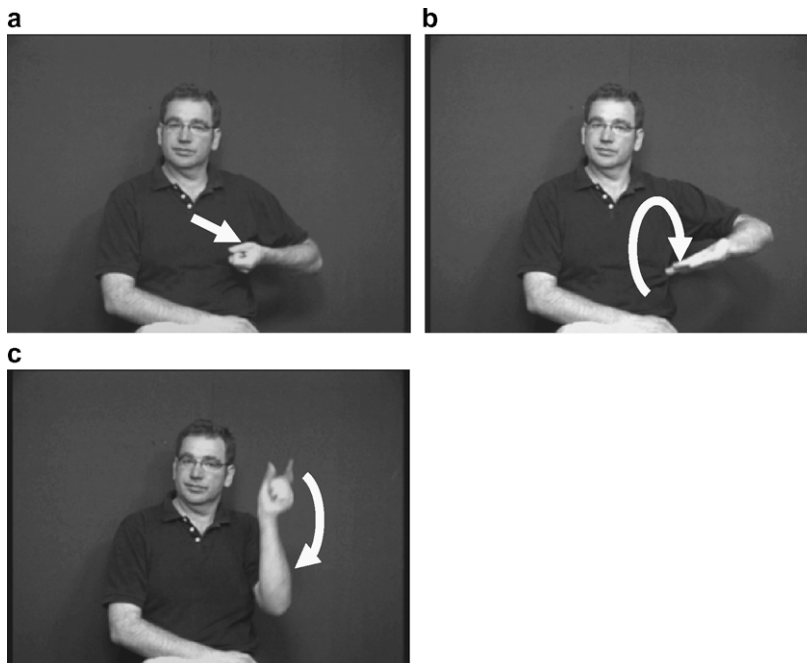
Fig. 2. An impossible nonsense sign and target in three video stills. At the start of the impossible nonsense sign the hand moves forward in a straight line (a). Then it moves upwards in an arc and simultaneously changes orientation and opens to form a new handshape (b). The target sign MOON which followed is shown in (c).

Videos with examples of stimuli are available at http://www.staff.city.ac.uk/g.morgan/sign_segmentation/clips. All experimental materials are available on request from the authors.

Four other native Deaf signers checked whether the possible signs were indeed nonexistent signs in BSL and its regional dialects. Any nonsense possible signs which any native signer said resembled a real sign were discarded.

They then judged the impossible signs and indicated those which they felt were problematic or marginally possible. We selected only those signs which all four native signers agreed were impossible.

In making the experimental stimulus pairs of possible and impossible nonsense we tried to match them for phonological complexity. This was achieved by including the same type (e.g., straight or arc movement) and number (e.g., one or two handshapes) of phonological parameters in the two contexts (handshape, path movement, internal movement, location, orientation and one or two hands). We rated nonsense signs as having 0 to 6 degrees of complexity (Sandler & Lillo-Martin, 2006). For example, if a nonsense sign had a marked handshape or two separate locations it was given one point of complexity. The average complexity of the nonsense signs in the nonsense combinations was 1.95. There was no difference between the complexity of the 1st and 2nd nonsense signs of each combination ($t(63) = -1.3$, $p = .198$). (See Orfanidou et al. (2009), for further details.)

We constructed 64 fillers consisting of two nonsense signs (32 combinations of two possible nonsense signs, and 32 combinations of an impossible nonsense sign followed by a possible nonsense sign). The second author practiced each sign in isolation and then produced them in the prescribed two-sign sequences. The materials were filmed in a professional studio; video clips were then edited into separate files using iMovie software.

There was a transitional movement between the nonsense sign and the context in 30 out of the 64 nonsense contexts (12 possible and 18 impossible contexts). Note that transitional movements and the lack of them are both common in sign languages (there can be no transition either because two consecutive signs are articulated in the same location, or because the locations are very similar so there is only a very small movement, or because, in rapid signing, there can be coarticulation; see Brentari (1998), for a full description of sign transitions). To our knowledge there is no study in BSL or any other sign language which has examined which pattern (transition or no transition) is the most frequent. For this reason we included both patterns. The duration and nature of the transition movements were similar in the two types of contexts. For example, it could be a transition from the nose to the forehead (for the possible context for a given target) paired with one from in front of the face to the forehead (for the impossible context), or from the shoulders to the neutral space (for both nonsense contexts for a given target). For the 32 targets, there were only three pairs of contexts in which there was a difference within the pair with regard to the presence or absence of a transition across the possible and impossible contexts (in two out of the three pairs the transition was present in the impossible context but not in the possible context). In these cases the difference between a transition and a lack of transition was kept as small as possible (e.g., a transition from eye to forehead paired with no transition at the forehead).

Experimental design

The target-bearing materials were split into two counterbalanced subsets, each with all 32 targets, but with 16 embedded in possible nonsense contexts and 16 in impossible nonsense contexts. Two lists were then created in which each subset of target-bearing items was mixed pseudo-randomly with the fillers (such that there was always at least one filler between two target-bearing items). List 1 included targets 1–16 with a possible nonsense context and targets 17–32 with an impossible nonsense context. The reverse was true for list 2. The only difference between the two lists concerned the contexts in which the targets appeared (i.e., the targets were in the same order in the two lists). Two additional lists were made with reverse orders of presentation. Each of the resulting four versions was split into two blocks of 48 trials (16 target-bearing items and 32 fillers).

Procedure

Each session started with a practice block (10 trials, modeled on the different types of trial in the main experiment), followed by the two blocks of the experiment proper. Stimuli were presented on a 19 in. computer screen using DMDX software (Forster & Forster, 2003). Each trial lasted 7 s, with 5.5 s for stimulus presentation and time for response, and then 1.5 s after time-out before the next trial began. The instructions to the participants were the following (translated from BSL): "you are going to watch two-sign sequences, and in some of them there will be a real sign. When you see a real sign press the button as fast as possible and then tell us what sign you saw". Through the practice items we showed participants that there were only two signs per trial, and that these were not classifier signs or compound signs. Participants all reported that they understood the requirements of looking for a real sign in sequences of two signs. Participants pressed the right button on a button-box if they spotted a real BSL sign embedded in the nonsense and then signed to a video camera in front of them what the sign was. RTs were measured from stimulus onset, but adjusted prior to analysis to measure from target offset (i.e., by subtracting video durations, as computed by multiplying the frame rate by the number of frames per stimulus; this proved to be the best estimate of how long the videos played during the experiment). Trials where participants pressed the button but then failed to sign the correct target to the camera (81; 4% of target-bearing trials) were treated as errors. There were also false alarms on fillers, that is, misperceptions of nonsense signs as real BSL signs. There were 678 false alarms (20% of the filler trials). Of those, 430 trials involved a misperception error. The remaining 248 trials were instances where participants thought they saw a sign (i.e., pressed the response button), but then copied exactly the nonsense sign they saw. For an analysis of false alarms, see Orfanidou et al. (2009).

Results and discussion

Mean Reaction Times (RTs, measured from video onset and offset) and mean errors for sign-spotting performance are shown in Table 1. Analyses of variance (ANOVAs) by participants (F1) and items (F2) were performed on both of the RT measures and on the errors, with Context (possible vs. impossible nonsense signs) as a within-participant

Table 1
Experiments 1 and 2: Mean Reaction Time (RT, in ms from target onset and offset) and mean error (proportions) in each context condition (Standard Errors in parentheses).

Experiment 1: sign spotting						Experiment 2: sign lexical decision					
Signs in possible-sign contexts			Signs in impossible-sign contexts			Signs taken from possible-sign contexts			Signs taken from impossible-sign contexts		
RT		Error	RT		Error	RT		Error	RT		Error
Onset	Offset		Onset	Offset		Onset	Offset		Onset	Offset	
2702 (10.0)	252 (38.8)	0.22 (0.02)	2742 (9.7)	309 (39.4)	0.25 (0.03)	1089 (64.5)	−67 (59.3)	0.14 (0.03)	1131 (60.8)	−31 (63.4)	0.17 (0.03)

and within-item factor and Version (4 counterbalanced versions of the materials) as a between-participant factor. RT analyses from target offset control for differences in stimulus (context and target) duration, and are the norm in the spoken variant of the word-spotting task (McQueen, 1996). But video duration can only be estimated, since actual playing times may vary across trials (J. Forster, pers. comm., June 26th 2009). We thus also report RT analyses from video onset.

Offset RTs

The effect of Context was statistically significant by participants only ($F1(1,56) = 11.9$, $p < .001$; $F2 < 1$; $min\ F'(1,34) = 0.467$, $p = .501$). To examine effects of AoA, additional by-participant ANOVAs were run with Group (3 levels: acquisition before age 5, between ages 6 and 12, or after age 12) as a second between-participant factor. There was no main effect of Group, and no interaction of Group with Context (all $F1$'s < 1), but the main effect of Context was still significant by participants (RTs, $F1(1,48) = 6.5$, $p < .015$, $F2 < 1$, $min\ F'(1,111) = 0.703$, $p = .403$). Participants were faster, on average, by 57 ms (±95%CI = 27 ms) in detecting real BSL signs in nonsense contexts that were possible BSL signs than in contexts that were impossible BSL signs.

Onset RTs

The effect of Context was again significant by participants only ($F1(1,56) = 6.3$, $p = .015$; $F2 < 1$; $min\ F'(1,33) = 0.188$, $p = .667$). In this analysis participants were, on average, 40 ms (±95%CI = 28 ms) faster to detect the target BSL signs in the possible BSL sign contexts. Additional ANOVAs were again run with Group as a second between-participant factor. There was no main effect of Group, and no interaction of Group with Context (all F's < 1), but the main effect of Context was significant by participants (RTs, $F1(1,48) = 6.5$, $p = .014$, $F2 < 1$, $min\ F'(1,91) = 0.04$, $p = .840$).

Errors

The Context effect in error rates was significant by participants and items ($F1(1,56) = 13.0$, $p < .001$; $F2(1,31) = 9.4$, $p < .004$; $min\ F'(1,32) = 5.5$, $p < .022$). Participants missed fewer target BSL signs (mean difference = 3%; ±95%CI = 4%) in nonsense contexts that were possible BSL signs than in contexts that were impossible BSL signs. In an analysis which included the factor Group, there was again no effect of Group or interaction of Group with Context (all F's < 1), but there was still an effect of Context ($F1(1,48) =$ 12.7, $p < .001$, $F2(1,93) = 4.6$, $p = .035$, $min\ F'(1,138) = 3.4$, $p = .068$).

In summary, Deaf signers of BSL were faster and more accurate spotting real BSL signs embedded in nonsense-sign contexts when the nonsense signs were possible BSL signs than when they were impossible BSL signs. But each nonsense sequence in Experiment 1 was produced separately. Hence the target signs in the impossible-sign contexts were physically different from those in the possible-sign contexts. Possible differences in fluency, clarity or speed of articulation of targets across contexts could thus account for the differences in sign-spotting performance across conditions. Experiment 2 was run to control for this possibility. As in many spoken-word-spotting studies (e.g., Cutler & Norris, 1988), the targets were digitally excised from their contexts and presented to new participants in a go/no-go lexical-decision task (i.e., participants had to press a button every time they saw a real BSL sign in a list of words and nonwords). We hoped to show that there would be no difference on lexical decision performance between signs taken from possible-word contexts and signs taken from impossible-word contexts.

Experiment 2

Method

Participants

Twenty native Deaf signers took part. None had participated in Experiment 1, and all had normal or corrected vision. They were paid to take part.

Stimuli, design and procedure

Each of the target signs from Experiment 1 was excised from its context using iMovie software. We took as the starting point for the target the point in time at which the handshape of the target had been formed. Fillers were created by excising, using the same criterion, the second nonsense sign in each of the Experiment 1 fillers. The experiment was exactly analogous to Experiment 1 (i.e., same critical materials, design, and running order) except that each target and each filler was presented without its original context. The instructions of Experiment 1 were modified slightly: Participants were asked to press the button whenever they saw a real BSL sign (they again signed targets that they detected to a video camera). Raw RTs were again adjusted to measure from the temporal offset of the target words. Button presses with incorrect repro-

ductions of the signs (17; 3% of target trials) were again counted as errors.

Results and discussion

Table 1 shows the mean RTs and mean error rates on lexical decisions in Experiment 2 to the Experiment 1 targets after the targets had been extracted from their contexts. In ANOVAs parallel to those in Experiment 1 there was, as predicted in this control experiment, no effect of the context from which the targets had been taken, either in offset RTs ($F1(1,16) = 1.1$, $p = .308$, $F2 < 1$, *min* $F'(1,36) = 0.078$, $p = .781$) or errors ($F1$, $F2 < 1$). The mean offset RT difference between possible and impossible context was 37 ms ($\pm 95\%CI = 82$ ms); the mean error difference was 3% ($\pm 95\%CI = 16\%$). In the analysis with onset RTs there was again no effect of context ($F1(1,16) = 1$, $p = .326$, $F2 < 1$, *min* $F'(1,47) = 0.330$, $p = .568$). The mean onset RT difference between possible and impossible contexts was 41 ms ($\pm 95\%CI = 75$ ms).

Experiment 2 shows that the differences in the ease of spotting real BSL signs when they were embedded in possible-sign relative to impossible-sign contexts were not due to differences in the way the real signs had been articulated in the different contexts. The comparison in Experiment 1 between impossible-sign contexts consisting of superfluous and illegal combinations of phonological parameters and possible-sign contexts with legal but non-existent parameter combinations raises an additional concern, however. Might the disadvantage for the impossible-sign context condition reflect a dispreference for contexts with illegal phonotactics rather than a dispreference for a lexically nonviable sign? That is, the disadvantage in the impossible-sign condition might not necessarily reflect the signs' status as impossible words. As we noted earlier, because of physical constraints, impossible BSL signs can only be made by creating phonotactically illegal signs. There thus appears to be no way to avoid this problem in a BSL experiment. In addition, it is not even known whether, in speech segmentation, a syllable containing phonotactically illegal material creates a segmentation problem. Since no previous word-spotting study on speech segmentation has made this test, Experiment 3 was designed. Dutch listeners were asked to spot spoken Dutch words. The target words were embedded in three nonsense contexts, two with syllables, and one with a single consonant (e.g., *long*, "lung", in *schruuslong*, *sfruuslong*, and *slong*). The first type of syllabic context had an onset consonant cluster that is legal in Dutch, the second had an illegal cluster.

We predicted that the listeners in Experiment 3 would find it harder to spot words in the single-consonant contexts (single consonants are not possible Dutch words) than in the contexts with syllables with legal clusters (which are possible Dutch words). This would replicate earlier results showing sensitivity to the PWC in Dutch (McQueen, 1998; McQueen & Cutler, 1998). Experiment 3 also tested the claim that the PWC in spoken-language is a simple vowel-based constraint (Cutler, Demuth, & McQueen, 2002; Norris, McQueen, Cutler, Butterfield, & Kearns, 2001; Norris et al., 1997): As implemented in the

Shortlist B model (Norris & McQueen, 2008), candidate words are dispreferred only when residues have no vowel. According to this view, word-spotting performance should be equally easy in the two syllabic contexts, since both contain vowels. That is, a syllable with a phonotactically illegal cluster would be considered to be as viable a residue as one with a legal cluster. One reason why this might be the case is that a syllable with an illegal cluster could be a foreign word or name. As Norris and McQueen (2008) argue, the current vowel-based implementation of the PWC in Shortlist B helps prevent recognition problems from arising when "unknown words" such as foreign names are encountered. In a Bayesian model such as Shortlist B, accepting such sequences as a potential part of the lexical parse means not only that the new word can be learned, but also that the words before and after it can be recognized.

The demonstration that word spotting in consonant contexts is harder than in syllabic contexts, and that there is no difference between the syllables with legal and illegal clusters, would thus confirm that the PWC in speech processing is a simple vowel-based constraint. More importantly for current purposes, however, this demonstration would also suggest that phonotactically illegal material, in keeping with the Shortlist B account, is in itself unlikely to create a speech segmentation problem. We could infer from this outcome of Experiment 3, on the grounds of parsimony, that the results of Experiment 1 are more likely to be because the impossible contexts were not viable BSL words than because they were phonotactically illegal. This inference is not unproblematic, because it depends on the assumption that speech and sign are processed in the same way. It is of course possible that sign and speech are segmented differently, such that phonotactic illegality matters for sign but not speech segmentation, and that there is a lexical-viability constraint only in spoken language processing. Nevertheless, a simpler account would be that sign and speech are processed alike, such that phonotactic illegality would not matter for segmentation in either modality, but that lexical viability would matter in both.

This argument from parsimony is critically dependent on there being consistent outcomes across Experiments 1 and 3. Spoken-word-spotting performance in the illegal syllable contexts could after all be found to be poorer than in the legal syllable contexts. This alternative outcome would challenge the current vowel-based implementation of the PWC (Norris & McQueen, 2008), and it would suggest that phonotactically illegal material in a nonsense context is dispreferred in segmentation. This in turn would undermine an account of the results of Experiment 1 in which sign segmentation is based on lexical-viability constraints.

Experiment 3

Method

Participants

The participants were 36 students (native speakers of Dutch from the MPI subject pool with no hearing disorders). They were paid to take part.

Stimuli

The targets were 45 monosyllabic Dutch words (nouns or adjectives; e.g., *long*, "lung") with Consonant–Vowel–Consonant (CVC) or CVCC structures. Each appeared in three contexts: a nonsense syllable with a triconsonantal onset cluster that is legal in Dutch (e.g., [sxrys], making the target-bearing item *schruuslong*), a nonsense syllable with an illegal triconsonantal onset cluster (e.g., [sfrys], in *sfruuslong*), and a single consonant (e.g., [s], in *slong*). The triplets for each target word were matched in the following way: The coda consonant of the two syllables was the one that was used as the consonant context, and the rimes of the two syllables were identical. Legal onset clusters included [spr], [str], [skr], [sxr], and [spl]; illegal clusters included [sfr], [snr], [swr], [sfl], [sxl], [spw], [stw], [spj], and [sxn]. The consonants in the single-consonant condition were chosen such that, when combined with the onset consonants of the target words, they formed legal onset clusters (e.g. [sl] in *slong*). All triplets are listed in Appendix A.

There were 96 fillers containing no Dutch words, modeled on the target-bearing items. There were 96 different CVC or CVCC nonsense syllables, preceded by three different nonsense contexts (32 with legal CCCVC nonsense syllables, 32 with illegal CCCVC nonsense syllables with legal onsets, and 32 with single consonants). All materials were digitally recorded five times by a native male speaker of Dutch in a sound-damped booth at 44.1 kHz and edited using Praat (Boersma & Weenik, 2005). Target words were excised from recordings of the single-consonant context items (e.g., *long* from *slong*), cutting the waveform at zero-crossings, and spliced onto contexts excised in a similar manner from each of the two syllabic contexts and from a different recording of the single-consonant context than that used as the source of the word. Target words were thus physically identical across the three contexts. (Note that a lexical-decision control for Experiment 3 was thus unnecessary.) Filler items were not spliced.

Experimental design

Following a similar counterbalancing procedure to Experiment 1, the target-bearing items were divided across three versions of the experiment, such that each version contained all 45 targets, 15 in each of the three contexts, and all 96 fillers, with targets and fillers in the same pseudo-random order in each version.

Procedure

A practice block (18 trials) was followed by one of the three versions of the main experiment. Stimuli were presented over Sennheiser headphones at a comfortable listening level in a sound-damped booth. Stimulus presentation, timing and data collection were controlled by NESU software (http://www.mpi.nl/world/tg/experiments/nesu.html). There were 4 s from stimulus onset on one trial to stimulus onset on the next trial. Participants were asked to press a button if they spotted a real Dutch word, and then to say aloud what they had spotted. RTs were recorded from stimulus onset, but adjusted prior to analysis so as to measure from target offset (by subtracting total stimulus duration from raw RTs), hence controlling for differences in context duration. Spoken responses were checked: Trials where the button was pressed but the target was not correctly named (72; 4.4% of target-bearing trials) were treated as errors.

Results and discussion

Mean RTs and mean error rates for spoken-word-spotting are presented in Table 2. ANOVAs by participants and items were performed on adjusted RTs and errors with Context (syllables with legal consonant clusters; syllables with illegal consonant clusters; and single consonants, i.e., impossible words) as within-participant and within-item factor and Version (3 counterbalanced versions of the materials) as a between-participant factor. There was a Context effect in RTs ($F1(2,66) = 146.2$, $p < .001$; $F2(2,88) = 60.6$, $p < .001$; *min* $F'(2,143) = 43$, $p < .001$) and errors ($F1(2,66) = 17.3$, $p < .001$; $F2(2,88) = 26.8$, $p < .001$; *min* $F'(2,136) = 11$, $p < .001$). Words in single-consonant contexts were spotted, on average, 390 ms more slowly (±95%CI = 62 ms) and 18% less accurately (±95%CI = 7%) than words in legal-cluster syllabic contexts (RTs: $F1(1,33) = 164.7$, $p < .001$, $F2(1,44) = 69.6$, $p < .001$, *min* $F'(1,72) = 48.9$, $p < .001$; Errors: $F1(1,33) = 27.4$, $p < .001$, $F2(1,44) = 47.3$, $p < .001$, *min* $F'(1,66) = 17.3$, $p < .001$). In addition, words in single-consonant contexts were spotted, on average, 382 ms more slowly (±95%CI = 53 ms) and 16% less accurately (±95%CI = 9%) than words in illegal-cluster syllabic contexts (RTs: $F1(1,33) = 211.3$, $p < .001$, $F2(1,44) = 96.0$, $p < .001$, *min* $F'(1,73) = 66$, $p < .001$; Errors: $F1(1,33) = 15.3$, $p < .001$, $F2(1,44) = 37.7$, $p < .001$, *min* $F'(1,58) = 10.8$, $p < .002$). In contrast, the differences between the syllables with legal and illegal clusters (mean RT difference = 10 ms, ±95%CI = 42 ms; mean error difference = 2%, ±95%CI = 5%) were not statistically significant (all F's < 1).

Experiment 3, using spoken Dutch, replicated the findings from BSL in Experiment 1. We suggest that the parallel effects across modalities are likely to be due to the same segmentation algorithm, even though there are differences in what can constitute an impossible word in sign versus speech. Experiment 3 also provided support for the current

Table 2
Experiment 3: Mean Reaction Times (RT, in ms from target offset) and mean errors (proportions) for the speech contexts (Standard Errors in parentheses).

Spoken words in syllable contexts				Spoken words in single-consonant (impossible-word) contexts	
Syllables with legal consonant clusters		Syllables with illegal consonant clusters			
RT	Error	RT	Error	RT	Error
808 (27.8)	0.31 (0.02)	816 (28.5)	0.33 (0.02)	1198 (24.6)	0.49 (0.04)

implementation of the PWC in the Shortlist B model (Norris & McQueen, 2008). In particular, it appears that syllables with illegal clusters are viable residues in speech segmentation. This allows listeners to recognize such sequences, should they occur in running speech (for example as foreign names), and to recognize the words surrounding those novel sequences.

General discussion

We tested here whether language segmentation processes in speech and sign differ, depending on the nature of the sensory input or, alternatively, whether they are the same, because of commonalities in the underlying computational problem. The results of Experiment 1 support the latter hypothesis. Signers of BSL found it harder to spot real BSL signs in nonsense contexts that were impossible BSL signs than in contexts which were possible BSL signs. Experiment 2 showed that this difference could not be due to physical differences in the way the target signs were articulated in the two types of context. The findings from Experiment 1 thus mirror previous results on the segmentation of spoken languages: Listeners find it harder to spot spoken words in impossible-word than in possible-word contexts. It appears that, despite the major differences between the sign and speech sensory signals, language segmentation proceeds so as to avoid leaving residues of the input that cannot themselves be lexical entries, as the PWC predicts.

Experiment 3 showed in addition that Dutch listeners find it as easy to spot words in the context of syllables with legal consonant clusters as in the context of syllables with illegal clusters; both of these contexts were easier than those with single-consonant clusters. The large difference between syllabic and consonantal contexts replicates previous results in Dutch (McQueen, 1998; McQueen & Cutler, 1998), and the lack of a difference between the two syllabic contexts confirms that the PWC in spoken language is a simple, vowel-based constraint (Norris & McQueen, 2008; Norris et al., 1997): Residues are only impossible words if they do not contain a vowel. The syllables with illegal clusters are of course not legal Dutch words, but the critical point here is that, despite the difference in the onset clusters, both syllable types contain vowels. Thus, with respect to lexical segmentation, it does not appear to matter whether a chunk of speech contains an illegal consonant cluster, so long as it also contains a vowel. This makes good sense: A residue without a vowel could never be a word that a speaker would intend, but a syllable with an illegal cluster could be a foreign word or name with different phonotactic restrictions.

The primary motivation for Experiment 3, however, was to address the concern that the difference found in Experiment 1 could be because of the presence of illegal combinations of parameters in the impossible-context condition (the impossible signs necessarily had superfluous parameters that the possible signs lacked) rather than because such signs cannot be BSL words. The results of Experiment 3 suggest that, at least in speech segmentation, the presence of phonotactically illegal material does not in itself create a segmentation problem. We therefore infer that the same is true in sign segmentation, and hence that the effect observed in Experiment 1 reflects differences between the contexts in their lexical viability rather than simple differences in the nature of the phonological material between contexts. Although this inference depends on a comparison between Dutch and BSL, and on the assumption that sign and speech are segmented in the same way, we favor it because it supports the most parsimonious theory of the results as a whole. The alternative account, that sign comprehenders are sensitive to phonotactic legality but speech comprehenders are not, and that sign comprehenders are not sensitive to lexical-viability constraints but speech comprehenders are, is more complex. The simplest account consistent with the present data is that, even though there are necessary differences in what can constitute an impossible word between speech and sign, the PWC applies in both cases. Further empirical support for this conclusion, preferably based on sign language experiments and hence requiring no assumption of equivalence across modalities, would nevertheless be welcome.

The present study adds to a growing body of work on sign recognition (Carreiras et al., 2008; Corina, 1993; Corina & Hildebrandt, 2002; Dye & Shih, 2006; Emmorey & Corina, 1990). It is the first, however, to examine the online sign segmentation process (though off-line segmentation judgments have been examined; Brentari, 2006). Effects of the density of the lexical neighborhood in Spanish Sign Language on sign recognition (Carreiras et al., 2008) suggest that there is a process of competitive evaluation of multiple sign hypotheses, akin to that in speech recognition. Although competition effects on sign segmentation have not been tested directly, it is thus plausible to assume that, as in speech segmentation, segmentation of sign is achieved through a competitive process. This proposal would fit well with the present results. In speech, the PWC is assumed to operate by modulating the lexical competition process (Norris & McQueen, 2008; Norris et al., 1997). The PWC in sign processing is likely to act in the same way: Sign hypotheses occurring in parses with impossible-sign residues would be less likely to win the competition.

Our findings also show that BSL comprehenders use knowledge based on the well-formedness rules of BSL. Several studies have shown that listeners use their knowledge of phonological structure constraints in segmenting speech (Dumay, Frauenfelder, & Content, 2002; McQueen, 1998; Suomi, McQueen, & Cutler, 1997; Warner, Kim, Davis, & Cutler, 2005; Weber & Cutler, 2006). Although we have argued that the results of Experiment 1 are not due to the phonotactic illegality of the impossible signs alone, but rather to the fact that these combinations are not viable BSL signs, these signs have that status because they break the language's well-formedness rules. Segmentation across languages thus appears to depend on phonotactic knowledge.

Evidence for the PWC has been found in many languages, including English (Norris et al., 1997, 2001), Dutch (McQueen, 1998; McQueen & Cutler, 1998), Japanese (McQueen, Otake, & Cutler, 2001), Sesotho (Cutler et al., 2002), Cantonese (Yip, 2004), German (Hanulíková,

2008), and Slovak (Hanulíková, McQueen, & Mitterer, in press). The lexical-viability constraint in the segmentation of these languages does not appear to vary as a function of language-specific phonological constraints on what can be a well-formed word. In all these cases, residues without vowels created segmentation problems, and those with vowels did not. Even in a language such as Slovak, which has single-consonant prepositions, word spotting is harder in contexts with nonprepositional consonants than in syllabic contexts (Hanulíková et al., in press). Although the consonants which can be prepositions in Slovak seem to be treated in Slovak segmentation as viable residues, this appears to be because they are themselves Slovak words. These consonants are treated specially by Slovak listeners, but the simple vowel-based constraint still applies to other consonants in Slovak. The PWC thus appears to apply across typologically diverse spoken languages. Critically, however, this prior work has ignored the existence of signed languages, and hence the fact that not all natural languages are spoken. The present results thus suggest that the PWC is not only a *speech* segmentation algorithm but instead that it has broader scope: The PWC appears to be a modality-general *language* segmentation algorithm.

Why should this be so? We suggest that a lexical-viability constraint operates in the segmentation of natural languages because of the benefits that it offers the perceiver. The PWC has been motivated because it helps the listener to segment continuous speech, to rule out spurious candidate words, and to deal with novel words in spoken input (Norris & McQueen, 2008; Norris et al., 1997). Parallel computational problems confront sign comprehenders who are faced with continuous signed input, and hence, we argue, they use the same solution to these problems. Signers appear to monitor the visual patterns in the sign stream and to compute when possible segmentations would leave lexically nonviable residues. Although in our experimental design we had to use impossible signs that would not occur in fluent BSL, we assume that the PWC would act during normal segmentation to penalize lexical parses with impossible residues (e.g., a parse leaving a movement parameter unattached to location and handshape parameters). This would help signers to rule out spurious lexical candidates during sign recognition, and to generate plausible hypotheses for new signs when novel input is encountered.

Finally, it is important to note that there was no AoA effect in Experiment 1. It would appear that the ability to use a lexical-viability constraint can be acquired even if there is no exposure to BSL until after the age of 12. This outcome is striking, given previous suggestions that native and nonnative signers process sign language in qualitatively different ways (Mayberry & Eichen, 1991). We suggest that this apparent paradox can be resolved simply by assuming that AoA does not affect all components of sign processing equally (see Emmorey et al. (1995) and Newport (1990) for similar suggestions). In the case of segmentation, it would appear that all sign comprehenders can benefit from a lexical-viability constraint, and hence that they all learn to reap this benefit, irrespective of AoA. The fact that sign perceivers who acquire BSL later in life can learn to use the PWC suggests that relatively lit-

tle exposure to language is required to support this learning. Indeed, 12-month-old infants learning American English already appear to use the PWC to segment speech (Johnson, Jusczyk, Cutler, & Norris, 2003). The lack of an AoA effect thus strengthens the argument about the modality-general utility of the PWC algorithm. The PWC provides a solution to the computational problem that continuous natural-language input poses to speakers and signers alike. Although the sensory inputs of sign and speech are very different indeed, their recognition depends on the same segmentation procedure.

Acknowledgments

Appendix A

Experiment 1: 32 targets (real BSL signs)

Dog, duck, nut, light, egg, wine, curtains, castle, moon, police, house, tree, monkey, trousers, shirt, banana, aeroplane, sister, name, piano, alarm, sport, weekend, afternoon, boy, theatre, university, gold, holiday, umbrella, hospital, bank.

Experiment 3: 45 targets (Dutch words), with contexts and translations

Target	Single-consonant context	Legal consonant-cluster syllable context	Illegal consonant-cluster syllable context	Target translation
rok	prok	sprooprok	sfrooprok	skirt
raaf	traaf	schrutraaf	sglutraaf	raven
rib	prib	spluuprib	sfluuprib	rib
roem	kroem	straukroem	spwaukroem	fame
reus	freus	skrafreus	spjafreus	giant
ramp	framp	skreuframp	snreuframp	disaster
lap	plap	struiplap	sgnuiplap	cloth
loods	floods	strufloods	sglufloods	shed
long	slong	schruuslong	sfruuslong	lung
lift	klift	splieklift	sflieklift	lift
luis	gluis	sproogluis	sfroogluis	louse
nut	snut	sprisnut	sgnisnut	use
worm	tworm	struutworm	sfruutworm	worm
wet	kwet	sproekwet	sgloekwet	law
wijk	twijk	schreutwijk	spjeutwijk	neighborhood
rund	grund	sproegrund	stwoegrund	cattle
reuk	treuk	skreutreuk	snreutreuk	smell
rijm	krijm	schrakrijm	sfrakrijm	rhyme
riem	griem	spruigriem	spwuigriem	belt
reep	preep	spruupreep	spwuupreep	bar

Appendix A (*continued*)

Target	Single-consonant context	Legal consonant-cluster syllable context	Illegal consonant-cluster syllable context	Target translation
ras	tras	skruutras	sfluutras	breed
luik	pluik	schrupluik	sglupluik	hatch
lijm	plijm	skreuplijm	sfleuplijm	glue
lus	flus	splaaflus	spwaaflus	loop
lot	klot	spruuklot	spjuuklot	fate
lek	glek	streglek	swreglek	leakage
lamp	plamp	schruiplamp	stwuiplamp	lamp
werk	twerk	sprautwerk	spjautwerk	work
wiel	swiel	schriswiel	swriswiel	wheel
warm	twarm	spluutwarm	stwuutwarm	warm
riet	kriet	spluukriet	swruukriet	reed
ram	fram	striefram	sgniefram	ram
rouw	prouw	schrooprouw	swrooprouw	mourning
ring	fring	spleufring	stweufring	ring
rots	frots	splaufrots	sglaufrots	rock
raar	praar	schruupraar	swruupraar	strange
lied	glied	skrooglied	spjooglied	song
leem	sleem	sprausleem	snrausleem	loam
lol	flol	streuflol	sgneuflol	fun
lijf	plijf	struplijf	sfluplijf	body
lont	plont	skruuplont	stwuuplont	fuse
licht	klicht	splaaklicht	spjaaklicht	light
naald	snaald	strisnaald	sglisnaald	needle
wol	swol	stretswol	snretswol	wool
wind	twind	skritwind	sgnitwind	wind

References

Boersma, P., & Weenik, D. (2005). Praat: Doing phonetics by computer (Version 4.3.14) [Computer program]. <http://www.praat.org/> Retrieved 26.05.05.

Brentari, D. (1998). *A prosodic model of sign language phonology*. Cambridge, MA: MIT Press.

Brentari, D. (2006). Effects of language modality on word segmentation: An experimental study of phonological factors in a sign language. In S. Anderson, L. Goldstein, & C. Best (Eds.). *Papers in laboratory phonology* (Vol. VIII, pp. 155–164). The Hague: Mouton de Gruyter.

Carreiras, M., Gutierrez-Sigut, E., Baquero, S., & Corina, D. (2008). Lexical processing in Spanish Sign Language. *Journal of Memory and Language, 58*, 100–122.

Cho, T., McQueen, J. M., & Cox, E. A. (2007). Prosodically driven phonetic detail in speech processing: The case of domain-initial strengthening in English. *Journal of Phonetics, 35*, 210–243.

Corina, D., & Emmorey, K. (1993). Lexical priming in American Sign Language. In *Poster presented at 34th annual meeting of the psychonomic society*. Washington, DC.

Corina, D., & Hildebrandt, U. (2002). Psycholinguistic investigations of phonological structure in American Sign Language. In R. P. Meier, K. Cormier, & D. Quinto-Pozos (Eds.), *Modality and structure in signed and spoken languages* (pp. 88–111). Cambridge: Cambridge University Press.

Corina, D., & Knapp, H. P. (2006). Lexical retrieval in American Sign Language production. In L. M. Goldstein, D. H. Whalen, & C. T. Best (Eds.), *Papers in laboratory phonology 8: Varieties of phonological competence* (pp. 213–240). Berlin: Mouton de Gruyter.

Cutler, A., Demuth, K., & McQueen, J. M. (2002). Universality versus language-specificity in listening to running speech. *Psychological Science, 13*, 258–262.

Cutler, A., & Norris, D. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology: Human Perception and Performance, 14*, 113–121.

Davis, M. H., Marslen-Wilson, W. D., & Gaskell, M. G. (2002). Leading up the lexical garden-path: Segmentation and ambiguity in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 28*, 218–244.

Dumay, N., Frauenfelder, U. H., & Content, A. (2002). The role of the syllable in lexical segmentation in French: Word-spotting data. *Brain and Language, 81*, 144–161.

Dye, M. W. G., & Shih, S. (2006). Phonological priming in British Sign Language. In L. M. Goldstein, D. H. Whalen, & C. T. Best (Eds.). *Papers in laboratory phonology* (Vol. 8, pp. 241–263). Berlin: Mouton de Gruyter.

Emmorey, K., Bellugi, U., Friederici, A., & Horn, P. (1995). Effects of age of acquisition on grammatical sensitivity: Evidence from on-line and off-line tasks. *Applied Psycholinguistics, 16*, 1–23.

Emmorey, K., & Corina, D. (1990). Lexical recognition in sign language: Effects of phonetic structure and phonology. *Perceptual and Motor Skills, 71*, 1227–1252.

Forster, K. L., & Forster, J. C. (2003). DMDX: A windows display program with millisecond accuracy. *Behavioral Research Methods Instruments Computers, 35*, 116–124.

Hanulíková, A. (2008). Lexical segmentation in Slovak and German. Unpublished Ph.D. Thesis, Humboldt University Berlin.

Hanulíková, A., McQueen, J. M., & Mitterer, H. (in press). Possible words and fixed stress in the segmentation of Slovak speech. *The Quarterly Journal of Experimental Psychology*.

Johnson, E. K., Jusczyk, P. W., Cutler, A., & Norris, D. (2003). Lexical viability constraints on speech segmentation by infants. *Cognitive Psychology, 46*, 31–63.

Luce, P. A., & Pisoni, D. B. (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing, 19*, 1–36.

Marr, D. (1982). *Vision*. San Francisco, CA: W.H. Freeman.

Mathur, G., & Best, C. (2007). Three experimental techniques for investigating sign language processing. In *Paper presented at the 20th annual CUNY conference on human sentence processing*. La Jolla, CA.

Mattys, S. L., White, L., & Melhorn, J. F. (2005). Integration of multiple speech segmentation cues: A hierarchical framework. *Journal of Experimental Psychology: General, 134*, 477–500.

Mayberry, R. I. (1994). The importance of childhood to language acquisition: Evidence from American Sign Language. In J. C. Goodman & H. C. Nusbaum (Eds.), *The development of speech perception* (pp. 57–90). Cambridge, MA: MIT Press.

Mayberry, R., & Eichen, J. (1991). The long-lasting advantages of learning sign language in childhood: Another look at the critical period for language acquisition. *Journal of Memory and Language, 30*, 486–512.

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*, 1–86.

McQueen, J. (1996). Word spotting. *Language and Cognitive Processes, 11*, 695–699.

McQueen, J. M. (1998). Segmentation of continuous speech using phonotactics. *Journal of Memory and Language, 39*, 21–46.

McQueen, J. M., & Cutler, A. (1998). Spotting (different kinds of) words in (different kinds of) context. In *Proceedings of the fifth international conference on spoken language processing* (Vol. 6, pp. 2791–2794).

McQueen, J. M., Norris, D., & Cutler, A. (1994). Competition in spoken word recognition: Spotting words in other words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 621–638.

McQueen, J. M., Otake, T., & Cutler, A. (2001). Rhythmic cues and possible-word constraints in Japanese speech segmentation. *Journal of Memory and Language, 45*, 103–132.

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science, 14*, 11–28.

Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition, 52*, 189–234.

Norris, D., & McQueen, J. M. (2008). Shortlist B: A Bayesian model of continuous speech recognition. *Psychological Review, 115*, 357–395.

Norris, D., McQueen, J. M., Cutler, A., & Butterfield, S. (1997). The possible-word constraint in the segmentation of continuous speech. *Cognitive Psychology, 34*, 191–243.

Norris, D., McQueen, J. M., Cutler, A., Butterfield, S., & Kearns, R. (2001). Language-universal constraints on speech segmentation. *Language and Cognitive Processes, 16*, 637–660.

Orfanidou, E., Adam, R., McQueen, J. M., & Morgan, G. (2009). Making sense of nonsense in British Sign Language (BSL): The contribution of different phonological parameters to sign recognition. *Memory & Cognition, 37*, 302–315.

Orfanidou, E., Adam, R., McQueen, J., & Morgan, G. (2008). Mind the gap: Are non-lexical transitions between signs a segmentation cue in British Sign Language? In *Paper presented at the 14th international*

conference on architectures and mechanisms of language processing, 4–6 September 2008. Cambridge, UK.

Sandler, W., & Lillo-Martin, D. (2006). *Sign language and linguistic universals*. Cambridge, MA: Cambridge University Press.

Stevens, K. N. (1998). *Acoustic phonetics*. Cambridge, MA: MIT Press.

Stokoe, W. (1960). Sign language structure: An outline of the visual communication systems of the American deaf. In A. John (Ed.), *Studies in linguistics: Occasional papers*. Buffalo: University of Buffalo.

Suomi, K., McQueen, J. M., & Cutler, A. (1997). Vowel harmony and speech segmentation in Finnish. *Journal of Memory and Language, 36*, 422–444.

Sutton-Spence, R., & Woll, B. (1999). *The linguistics of BSL: An introduction*. Cambridge, UK: Cambridge University Press.

Uyechi, L. (1996). *The geometry of visual phonology*. Stanford, CA: CSLI Publications [1994].

Warner, N., Kim, J., Davis, C., & Cutler, A. (2005). Use of complex phonological patterns in speech processing: Evidence from Korean. *Journal of Linguistics, 41*, 353–387.

Weber, A., & Cutler, A. (2006). First-language phonotactics in second-language listening. *Journal of the Acoustical Society of America, 119*, 597–607.

Yip, M. C. W. (2004). Possible-word constraints in Cantonese speech segmentation. *Journal of Psycholinguistic Research, 33*, 165–173.