

Original Articles

Pre-linguistic segmentation of speech into syllable-like units

Okko Räsänen^{a,*}, Gabriel Doyle^b, Michael C. Frank^b^a Department of Signal Processing and Acoustics, Aalto University, P.O. Box 12000, Aalto, Finland^b Department of Psychology, Stanford University, Stanford, CA 94305, United States

ARTICLE INFO

Keywords:

Early language acquisition
Speech perception
Syllables
Speech segmentation
Oscillatory entrainment
Sonority

ABSTRACT

Syllables are often considered to be central to infant and adult speech perception. Many theories and behavioral studies on early language acquisition are also based on syllable-level representations of spoken language. There is little clarity, however, on what sort of pre-linguistic “syllable” would actually be accessible to an infant with no phonological or lexical knowledge. Anchored by the notion that syllables are organized around particularly sonorous (audible) speech sounds, the present study investigates the feasibility of speech segmentation into syllable-like chunks without any a priori linguistic knowledge. We first operationalize sonority as a measurable property of the acoustic input, and then use sonority variation across time, or speech rhythm, as the basis for segmentation. The entire process from acoustic input to chunks of syllable-like acoustic segments is implemented as a computational model inspired by the oscillatory entrainment of the brain to speech rhythm. We analyze the output of the segmentation process in three different languages, showing that the sonority fluctuation in speech is highly informative of syllable and word boundaries in all three cases without any language-specific tuning of the model. These findings support the widely held assumption that syllable-like structure is accessible to infants even when they are only beginning to learn the properties of their native language.

1. Introduction

Theories of early language acquisition often assume that infants perceive speech in terms of syllabic units, even before they can extract the words of their native language. For instance, many artificial language learning experiments have been conducted using stimuli whose statistics are manipulated at the syllabic level and where the success in word learning is measured in terms of the learner’s ability to capture statistical regularities connecting adjacent (Saffran, Aslin, & Newport, 1996) or non-adjacent (Newport & Aslin, 2004) syllables. Similarly, studies on artificial grammar learning have often used syllables as the representational level upon which the grammar operates (e.g., Gomez & Gerken, 1999; Marcus, Vijayan, Bandi Rao, & Vishton, 1999). The authors of these early behavioral studies were careful not to specify any specific type of representation underlying the statistical or rule-like computations capturing the syllable-level manipulations, simply referring to “statistical cues” and “speech sounds” (e.g., Saffran, Johnson, Aslin, & Newport, 1999; Saffran et al., 1996; Thiessen & Saffran, 2003), or “rules” and “variables” (Marcus et al., 1999). Nevertheless, later research has adopted the concept of syllable as a representational unit for pre-linguistic speech more explicitly (e.g., Frank, Goldwater, Griffiths, & Tenenbaum, 2010; Gambell & Yang, 2006; Meylan,

Kurumada, Börschinger, Johnson, & Frank, 2012; Perruchet & Tillman, 2010; Perruchet & Vinter, 1998; Swingley, 2005; see also Swingley, 2005, for a related discussion). For example, both SARAH (Mehler, Dupoux, & Segui, 1990) and WRAPSA (Jusczyk, 1993; see also Jusczyk & Luce, 2002) models of early speech perception assume that infants are capable of segmenting speech into syllable-like segments before further phonological and lexical analysis. Many Bayesian models of word segmentation also assume that syllable boundaries and identities are known as precursors to word recognition (Doyle & Levy, 2013; Phillips & Pearl, 2012).

The general approach of assuming syllables is consistent with empirical findings suggesting that infant speech perception is better characterized in terms of syllabic frames than phonemic segments (Bertoncini, Bijeljac-Babic, Jusczyk, Kennedy, & Mehler, 1988; Jusczyk, Bertoncini, Bijeljac-Babic, Kennedy, & Mehler, 1990; Jusczyk & Derrah, 1987; Jusczyk, Kennedy, & Jusczyk, 1995) and of holistic rather than analytic representations (Dupoux, 1993; see Hallé and Christia (2012), for an overview).¹ But despite this belief in the importance of syllables in language acquisition, adult-like syllabification depends on knowledge of phonological structure and of the specific language being used, neither of which is available to a child in the early stages of language acquisition. Existing developmental research has not been clear on what

* Corresponding author at: Department of Signal Processing and Acoustics, Aalto University, P.O. Box 12000, FI-00076 Aalto, Finland.

E-mail address: okko.rasanen@aalto.fi (O. Räsänen).

¹ Some research has even gone further and suggested that syllables remain the core unit in adults’ representations (Nasukawa, 2007; van der Hulst, 2005).

form of “syllabic units” would be available to pre-linguistic infants, nor how they could be identified at an age where children only have their generic perceptual capabilities to bootstrap their language learning. Is there language-independent information potentially available to a child in the speech signal itself that would permit the extraction of syllables or syllable-like objects as an early step in language learning?

Since young infants have not yet mastered the sound system of their native language, it is unlikely that the syllable-like chunks they perceive – which are our current focus here – would correspond precisely to the phonological syllables of the language in question. Phonological syllables are defined at the level of formal linguistic representations and typically include a number of language-specific rules and constraints (e.g., Clements, 1990; Goldsmith, 2011; Redford & Randall, 2005; Tesar & Smolensky, 1998), calling for significant experience with the language in question. Further complicating matters, there is some debate about the definition of a syllable among phoneticians dealing with speech perception and production in the wild. This debate is a product of the observation that phonological definitions do not always have clear counterparts in the measurable acoustic or auditory structure of speech available to the listeners (Ladefoged, 2000; Malmberg, 1963; Ohala, 1990; Palmer, 1978), and boundaries of phonological and phonetic syllables may differ in certain situations (e.g., French mute *e*; Fudge, 1969). Hence, the “syllable” is not a unanimous concept in the context of speech perception, even in the case of linguistically proficient adults (Price, 1980). Given this complex situation, the perceptually-available units we focus on here are not traditional phonological syllables, and therefore we will refer to them as *acoustic chunks*.

We explore the idea that these prelinguistic acoustic chunks are derived from sonority, the relative audibility of speech sounds. Sonority connects phonological and phonetic characterizations of syllables and – critically for our purposes – also is available as a possible cue for the creation of acoustic chunks for prelinguistic infants. All definitions of syllables, regardless of level, agree that they are related to the rhythmic fluctuation of speech sonority. A syllable minimally consists of a local sonority maximum (the nucleus), which is typically a vowel, and optionally of less sonorous sounds in the onset and the coda. Sonority also (typically) decreases monotonically from the nucleus towards the edges of the syllable (Hooper, 1976). Researchers may disagree on whether sonority can be defined in terms of physical or perceptual properties of speech (Clements, 2009; Galves, Garcia, Duarte, & Galves, 2002; Parker, 2002), whether it is a purely phonological abstraction based on structural description of languages similarly to phonemes (see Clements (1990), for a discussion), and whether sonority is relevant to phonological theories of language (Harris, 2006).

Irrespective of this debate, sonority has a strong correlational relationship with measurable properties of speech such as intensity and voicing (Section 1.2), and it was originally proposed as the perceptual audibility of different speech sounds (e.g., de Saussure, 1916; Jespersen, 1920; Whitney, 1874). Several existing automatic algorithms for speech syllabification are based on representations of acoustic speech that are closely related to sonority, such as low-pass filtered amplitude or energy envelopes (Mermelstein, 1975; Obin, Lamare, & Roebel, 2013; Wang & Narayanan, 2007). Correlates of sonority are also used in a number of recent neurophysiological models of speech perception, where the neural oscillations of the auditory cortex are believed to phase-lock to the rhythmic properties of the speech envelope, providing timing for more detailed analysis of speech sounds within the resulting syllabic frames (Ghitza, 2011; Giraud & Poeppel, 2012; Hyafil, Fontolan, Kabdebon, Gutkin, & Giraud, 2015). Even without a unanimous and physically precise definition for sonority (if such definition can ever exist), we can still investigate how syllabic structure is represented by rhythmic properties of speech by tracking measurable correlates of sonority as a property of the acoustic signal.

Given this background, in the present paper we examine the nature of the *acoustic chunks* that are accessible through rhythmic sonority variation in speech. Our rationale is that sonority information is likely

available to pre-linguistic infants, who are known to be sensitive to rhythmic properties of speech already at a young age (Jusczyk & Thompson, 1978; Mehler & Christophe, 1995; Mehler et al., 1988; Nazzi, Bertoncini, & Mehler, 1998), and hence sonority might be a cue that would allow infants to bootstrap acoustic chunks as inputs to other learning mechanisms. The method for our investigation is a computational model that instantiates this proposal of sonority-based segmentation into acoustic chunks. Before testing this proposal, we will motivate and clarify our goals by reviewing the evidence on the role of syllables in speech perception and on sonority in syllabic structure.

1.1. Role of syllables in speech perception

The idea that syllables² are central to pre-linguistic speech perception can be traced back to pioneering work by Mehler, Bertoncini, Jusczyk and their colleagues. In their seminal work, Bertoncini and Mehler (1981) showed that 2-month-old infants are better at discriminating acoustic differences between sequences that resemble syllables than sequences that do not, suggesting the role of syllable as a natural unit of speech perception. This result was later supported by similar findings with 4-day-old (Bijeljac-Babic, Bertoncini, & Mehler, 1993) and 3–4-month-old infants (Eimas, 1999). In addition, a series of studies revealed that 2-month-old infants are capable of extracting and retaining acoustic properties of syllables, and that the results are better understood in terms of syllabic, not phonemic, segmental units (Bertoncini et al., 1988; Jusczyk & Derrah, 1987; Jusczyk et al., 1990; Jusczyk et al., 1995). In parallel, a number of findings revealed perceptual primacy of syllables over phonemes in adult listeners (Mehler, Dommergues, Frauenfelder, & Segui, 1981; Segui, Frauenfelder, & Mehler, 1981; see also Mehler (1981), for a discussion). Studies with children (Lieberman, Shankweiler, Fischer, & Carter, 1974) and illiterate adults (Morais, Bertelson, & Alegria, 1986; Morais, Content, Cary, Mehler, & Segui, 1989) also show that conscious access and mental manipulation of syllables is easier than access to phonological segments for participants who have not received formal language instruction. In general, speech perception research shows that phone perception is conditioned by the neighboring sound context, indicating that temporal units of perception must be greater than individual phones (see, e.g., Nusbaum and DeGroot (1991), for a discussion). It is also well-known that syllables often exhibit coarticulatory patterns where cues for different speech sounds overlap in time (e.g., plosive cues are located in the formant transitions of the following vowel in CV-syllables).

Psychoacoustic and neurophysiological data also support the idea of perceptual processing of speech at time-scales greater than individual phone segments. The low-level auditory system integrates signal information across durations approximately corresponding to syllable lengths (~250 ms) (see, e.g., Wagner, 2008, or Räsänen & Laine, 2013, for a review) and is most sensitive to amplitude modulations around 4–5 Hz (e.g., Dau, Kollmeier, & Kohlraus, 1997; Viemeister, 1979)—a typical syllable rate in continuous speech (Greenberg, Carvey, Hitchcock, & Chang, 2003). Oscillatory neural activity in the auditory cortices is known to phase-lock to the syllable-driven amplitude envelope of the incoming speech input (Gross et al., 2013; Luo & Poeppel, 2007). The strength of the coupling correlates with the general intelligibility of the stimuli (Gross et al., 2013), as well as one’s subjective comprehension (Ahissar et al., 2001; Peelle, Gross, & Davis, 2013), with abnormalities in amplitude modulation tracking being associated with dyslexia or otherwise impaired phonological development (Leong & Goswami, 2014, 2015). These findings have led to neurophysiological models of speech perception where syllabic rhythm, cued by temporal modulations in the speech envelope and manifested as entrained theta-

² Here and below, we continue to use the term “syllable” to describe the units studied in prior research, even though we are – as noted above – agnostic about whether these are true syllables or merely acoustic chunks.

range oscillations in the brain, acts as a temporal chunking mechanism that provides frames of analysis for more detailed spectrotemporal analysis of the speech input (Ghitza, 2011; Giraud & Poeppel, 2012; Hyafil et al., 2015).

Lastly, the ability to extract syllables would also provide a great springboard for a language learner to infer word boundaries. A significant proportion of word tokens in conversational English speech are monosyllabic (e.g., 80% of tokens in conversational speech; Greenberg, 1999), and therefore successful syllabification of English would already lead to the correct segmentation of most words occurring in running speech (see also Gambell & Yang, 2006). The situation changes slightly for some other languages. French, Spanish, and Finnish, for instance, all have very different syllabic structure from English, with a much smaller set of unique syllables and much larger proportion of multisyllabic words in their vocabularies. Still, access to the syllabic structure of speech input would greatly limit the potential locations for word onsets and offsets even in these languages and provide useful temporal chunking of the input into perceptual units upon which further perceptual processing and learning can take place.

1.2. Sonority and syllables

What property of speech gives rise to the impression of syllable-like units? The existing attempts to define syllables on a physical basis can be divided into articulatory (e.g., Krakow, 1999; Stetson, 1951) and perceptual theories (developed by, e.g., de Saussure, 1916; Jespersen, 1920; Whitney, 1874; see also Malmberg (1963) and Goldsmith (2011), for an overview). In the present study, we will focus solely on perceptual information, following other research on early speech perception in prelinguistic infants.³

Considering the role of syllables in infant speech perception, Bertoncini and Mehler (1981) originally hypothesized that: “We consider continuous speech as being processed in segments or chunks of a few hundred milliseconds each. ... The delimitation of the chunks may be obtained by operating upon the maxima and minima in the wave envelope of the speech flow signal.” In other words, they suggested that the auditory system could use amplitude modulations of the speech envelope for chunking of the input into syllable-like segments. Even though they did not explicitly mention sonority, the conceptual idea of Bertoncini and Mehler is more generally known as the sonority sequencing principle (SSP; see Parker (2002), for an overview). Fig. 1 shows a simplified example of the principle on a segmental level where three monosyllabic words all have one highly sonorous vowel as the syllable nucleus, flanked by less sonorous consonants in the onset and the coda.

There is a broad phonological literature on sonority, covering its role in phonological theories, how different segments should be ranked in sonority, and how to account for language-specific properties for the SSP (e.g., Clements, 1990; Guffey, 2002). However, for the purpose of investigating infant speech perception—a case where phonemic discrimination skills have not yet developed and are likely to parallel lexical learning (see, e.g., Räsänen and Rasilo (2015), for a recent overview)—sonority has to be operationalized as a measurable property of speech independent of any phonological representations and constraints. This operationalization can be done either using straightforward physical correlates of sonority such as amplitude or intensity, or in terms of computational models of speech perception that use a more complex set of transformations that try to simulate how human hearing detects sonority fluctuations.

In fact, there are several existing attempts to measure syllabic rhythm from speech that have followed older theories of sonority as speech sound audibility (de Saussure, 1916; Jespersen, 1920; Whitney,




Fig. 1. A schematic example of the sonority sequencing principle at a segmental level, where each consonant (C) and vowel (V) segment is assigned with a relative sonority value. Syllables can be observed as units that have maximum sonority in the nuclei (typically V) and decreasing sonority towards the onset and offset. In this example, each syllable also makes up an individual word.

1874). Ohala (1992) argues that phonological definitions of sonority are circular, and proposes replacing the concept with a set of acoustic features including signal intensity and fundamental frequency. Parker (2002) considers five possible acoustic markers of sonority and argues sonority is best captured by signal intensity, which physically grounds the concept of sonority in speech envelope information. Jany, Gordon, Nash, and Takara (2007) follow Parker’s suggestion and find fairly consistent cross-linguistic sonority hierarchies based on intensity while Galves et al. (2002) have characterized sonority in terms of variation in spectral energy in a local temporal context. In his later works, Clements (2009) suggested that sonority is best understood in terms of perceptual “resonance” related to formant structure, although he did not explicate this further.

Other analyses related to deriving sonority-like measures for speech rhythm include a study by Leong, Stone, Turner, and Goswami (2014) who used an amplitude modulation (AM) filterbank to extract different modulation rates from speech, including the so-called syllable AM ranging in the frequency range from 2.3 to 7 Hz (see also Leong & Goswami, 2015), a study by Lee and Todd (2004) where a model of auditory periphery was used to measure rhythmic variations of prominence in speech, and a recent neurophysiological model by Hyafil et al. (2015) where spiking patterns of a set of interconnected excitatory and inhibitory neurons in the auditory cortex are entrained to the input from a model of the sub-cortical auditory system. In the field of speech technology, a number of algorithms have been proposed for automatic syllabification of speech that also use similar temporal modulations of speech energy to determine syllable boundaries as minima in the temporal envelope (e.g., Mermelstein, 1975; Obin et al., 2013; Villing, Timoney, Ward, & Costello, 2004) or speech rate as the number of sonority maxima (“syllable nuclei”) per time unit (Morgan & Fosler-Lussier, 1998; Wang & Narayanan, 2007).

1.3. This paper

In general, all the above approaches can be viewed as practical attempts to quantify rhythmic properties of speech, by explicitly or implicitly using some conceptualization of sonority as the alternation of relative prominence (or energy) between syllable nuclei and less audible sounds closer to syllabic boundaries. In order to play a functional role in the speech chain, the temporal sonority variation needs to have measurable properties in the acoustic speech, has to be coded by the auditory system, and is necessarily reflected in the brain responses and impacts behavioral responses to speech stimulation. These acoustic, perceptual, and neurophysiological factors can be seen as different viewpoints to the SSP-like organization of speech sounds according to their relative magnitude on some scale of measurement varying at typical syllabic rates, even if we do not know in detail what is the best way to define such scales in each of the representational domains or how they map to linguistic theories concerned with syllables.

In the present paper, we build on this background and describe a new oscillator-based model for sonority estimation and rhythmic segmentation that tries combine the best of the several worlds: applicability to real speech with modest computational resources, general

³ It is worth keeping in mind that articulatory development largely parallels early word learning (albeit with some delay) and the role of the possible interactions between speech perception and production in infancy are still largely unknown.

compatibility with the basic principles of the early auditory system and neurophysiological models of speech perception, and simplicity for the sake of model transparency and generality. Since purely bottom-up sonority-based segmentation without linguistic knowledge is likely to differ from the phonological syllabic structure of the language (see Appendix A for further discussion), we analyze how well acoustic chunks resulting from our model compare with the syllabification performance of three existing algorithms in conversational speech in three different languages (American English, Finnish, and Estonian), including both adult- and infant-directed speech in American English. Only limited cross-linguistic comparisons of segmentation accuracy have previously been carried out, and the present study allows us to ask 1) whether rhythmic cues are useful for the detection of syllables or words across these languages and speaking styles, and 2) how well a generic sonority-based segmentation system can perform in languages with different rhythmic characteristics before language-specific learning takes place.

2. An oscillator model for sonority-based rhythmic segmentation

In order to investigate rhythmic cues to syllabic structure in speech, we created a simple computational model of rhythmic perception that attempts to simulate human perception of sonority in speech.⁴ Following the principles of the recent neurophysiological models of speech perception (Ghitza, 2011; Giraud & Poeppel, 2012; Hyafil et al., 2015), the essence of the model is a model of oscillatory behavior that is controlled by the incoming speech input, that is, by using a set of damped oscillators that entrain to the amplitude modulations in speech. This provides a link between raw acoustic input arriving to the ear with the speech parsing mechanisms hypothesized to operate at the level of the cortex, combining the concepts of syllable-rate temporal modulations, sonority, oscillatory entrainment, and chunking of speech into syllable-like segments into a single unified and functional framework.

Consistent with the principles of psychoacoustic models of loudness perception, the model assumes that the overall perceptual sonority is a non-linear combination of amplitude modulation analysis at multiple frequency bands: Signals having strong and coherent excitation (phase) across multiple critical bands of hearing will lead to strong sonority perception (e.g., formants of voiced sounds; cf. also Clements, 2009, idea of “resonance”) whereas low overall energy (stops) or out-of-phase excitation will be perceived as less sonorant (see also Leong, 2012; Morgan & Fosler-Lussier, 1998; Wang & Narayanan, 2007). To implement this assumption, the model contains an auditory front-end that analyzes incoming speech on a number of frequency bands whose tuning curves follow what is known about the human auditory system. The amplitude envelope on each of the bands is then used to drive a damped harmonic oscillator (i.e., a modulation filter) that is tuned to the temporal modulations, or “rhythm”, of the input. A separate oscillator is used on each frequency band. The overall sonority envelope is obtained by combining the output from the set of band-specific oscillators, thereby forming a holistic representation of sonority fluctuation over time. Finally, “syllable nuclei” (or their corresponding realization within acoustic chunks) can be detected as maxima in the sonority envelope whereas “syllable boundaries” are defined as local minima between nuclei (see Figs. 2 and 3 below).

We make a simplifying assumption that each acoustic chunk spans the time interval between two sequential minima in the sonority envelope. This assumption allows straightforward comparison of the model output with a manually annotated linguistic reference that is also based on non-overlapping sequential units. It is likely, however, that rhythmic cues do not enforce a hard segmentation into temporally disjoint units but instead act as probabilistic and violable constraints in

the perceptual parsing of speech (cf., Prince & Smolensky, 1993). In case of hard segmentation, any errors in the front-end processing would propagate through the system and lead to a sub-optimal sequential representation of speech before the meaning of the message is even taken into account (see Räsänen & Rasilo, 2015). Investigating probabilistic cue combination in syllabic segmentation or questioning the entire existence of an explicit segmentation mechanism is beyond the scope of the present study. Instead, our present aim is to investigate the usefulness of sonority alone.

2.1. Technical details

Fig. 2 shows a schematic view of the model and Fig. 3 shows an example of the signal at different stages of the processing.

In order to simulate the frequency selectivity of the human ear, the incoming speech waveform (16,000 Hz) is first filtered into 20 logarithmically spaced frequency bands (100–7500 Hz) using a Gammatone filterbank⁵ (Patterson et al., 1992) (Fig. 3, second panel). The resulting amplitude envelope $e_c(t)$ on each frequency band $c \in \{1, 2, \dots, 20\}$ is then downsampled to have a sampling rate of $f_s = 1000$ Hz. In order to capture syllable-rate rhythmic variation in the amplitude envelopes, each envelope is fed to a harmonic oscillator with a center frequency f_0 and a given bandwidth Δf . The former corresponds to the frequency to which the oscillator is the most sensitive and the latter describes the range of input frequencies around the center frequency that are still reflected in the oscillator output. Both parameters are tuned to the typical range of syllabic structure (e.g., $f_0 = 5$ Hz and $\Delta f = 6$ Hz leads to a frequency band of $f_0 \pm \Delta f/2 = 2$ –8 Hz), and the same parameters are used for all band-specific oscillators.

A discrete-time oscillator driven by the envelope $e_c(t)$ can be described using the following equations:

$$\begin{aligned} F_c(t) &= e_c(t) - kx_c(t-1) - dv_c(t-1) \\ v_c(t) &= v_c(t-1) + F_c(t)/(f_s m) \\ x_c(t) &= x_c(t-1) + v_c(t)/f_s \end{aligned} \quad (1)$$

where $F_c(t)$, $v_c(t)$ and $x_c(t)$ are the force, velocity, and amplitude of the oscillator at band c at time t , respectively. Technically, this type of mechanical harmonic damped oscillator model corresponds to a second order electronic resonator (Oppenheim, Willsky, & Young, 1983) with its center frequency f_0 and ± 3 dB bandwidth determined by the mass m , spring constant k , and damping coefficient d of the oscillator. By fixing the spring constant to $k = 1$, the mass and damping coefficient can be used to control the center frequency f_0 and bandwidth Δf of the oscillator, respectively. In this study, oscillators for all bands had the same tuning.

Given the oscillator amplitudes $x_c(t)$ on each frequency band c at time t (Eq. (1)), the N most energetic bands at each time frame are combined into one sonority envelope $S(t)$ by taking the product of the amplitudes (see Wang and Narayanan (2007), for a similar procedure in speech-rate estimation). Since we are only interested in the relative instead of absolute sonority values over time, a numerically stable and computationally simpler method resulting in equivalent maxima and minima locations is to take a logarithmic sum of the amplitudes:

$$S(t) = \sum_{i=1}^N \log_{10}\{\vec{x}_i(t)\} \quad (2)$$

where $\vec{x}_i(t)$ is the set of oscillator amplitudes at time t sorted into descending order based on their amplitude values. Before applying the logarithm in Eq. (2), a fixed offset value ($-\min\{x_i\} + \varepsilon$) is added to all x_i to ensure that the oscillator amplitude is always positive ($x_i(t) \geq 0 + \varepsilon \forall t, i$) and output of Eq. (2) is therefore a real number. In

⁴ A MATLAB implementation of the oscillator model is available for download at <https://github.com/orasanen/thetaOscillator>

⁵ We use the Matlab implementation of Gammatone bank provided by Ning Ma available at <http://staffwww.dcs.shef.ac.uk/people/N.Ma/resources/gammatone/> (last accessed 23 November 2015).




Fig. 2. A block schematic of the model.

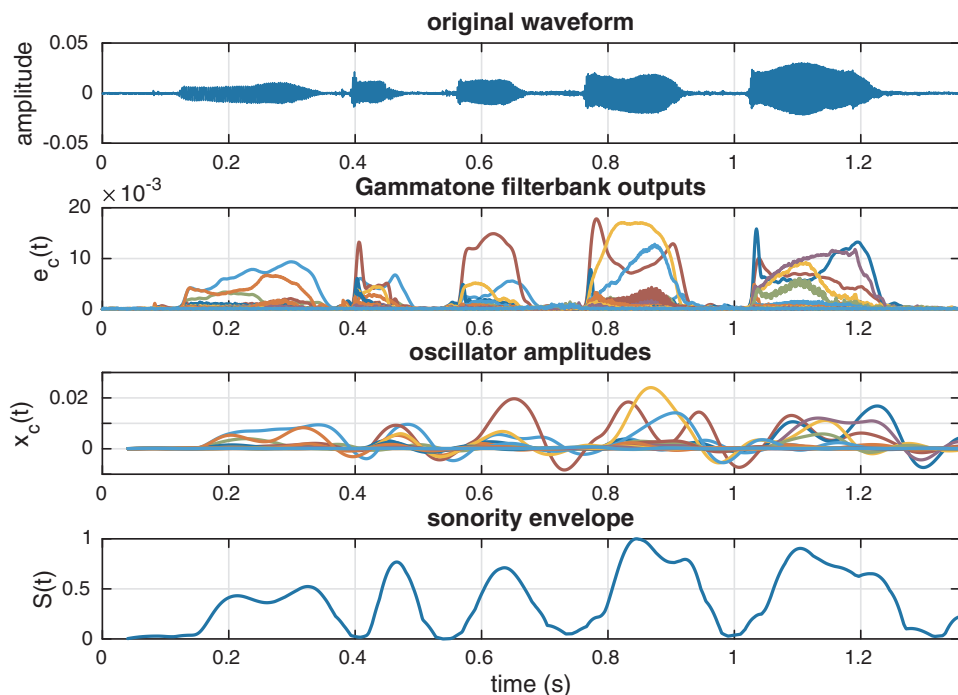


Fig. 3. An example of the processing stages for an utterance “Can you get the circle?”. Top panel: the original acoustic waveform. Second panel: the outputs $e_c(t)$ from the Gammatone filterbank, each frequency band c shown with a separate colored curve. Third panel: corresponding oscillator amplitudes $x_c(t)$. Bottom panel: the resulting sonority envelope $S(t)$ combined across band-specific amplitudes.

practice, Eq. (2) is analogous to calculating the overall cross-correlation between the frequency bands, but only considering the N most energetic bands⁶ out of all 20 bands at each moment of time. Preliminary testing revealed that N in the range of 6–16 leads to similar results and we decided to fix $N = 8$ for all experiments. Finally, in order to detect syllable boundaries, the sonority envelope S of the utterance is normalized to a value range of $[0, 1]$ and then each local minimum that is preceded by a local maximum of at least δ units higher is marked as a boundary between acoustic chunks.

3. Experiments

3.1. Data

The feasibility of the rhythm-based syllabic segmentation was evaluated using three corpora that were publicly available with manually verified syllable-level annotations: the Switchboard corpus, consisting of spontaneous telephone conversations in American English (Godfrey, Holliman, & McDaniel, 1992), the Phonetic Corpus of Estonian Spontaneous Speech (“EstPhon”) (Lippus, Tuisk, Salveste, & Teras, 2013), and the FinDialogue Corpus of Spontaneous Finnish Speech (Lennes, 2009). All three corpora consist of spontaneous speech between two conversing test subjects whose speech is recorded through telephone lines (Switchboard) or using near-field microphones (EstPhon and FinDiag).

These languages cover a range of syllable patterns. Unlike English, which is stress-timed, Finnish is considered to be syllable-timed and Estonian exhibits mixed timing patterns with elements from both stress- and syllable-timed languages (Grabe & Low, 2002). Similarly to

English, however, the primary stress is almost always on the first syllable of a word in Estonian and Finnish. In the Finnish syllable structure, CV is by far the most frequent type, CV(V) syllables covering 53.1% of all syllable tokens and 91% of tokens having a consonant onset. In contrast, consonant clusters (two or more subsequent consonants) are infrequent and occur mostly in loan words (for instance, CV(V)-syllables are estimated to be 45 times more frequent than (C) VCCs; see Suomi, Toivanen, and Ylitalo (2008), and references therein). Simple syllable structure also means that Finnish and Estonian words are almost always multisyllabic, whereas English has many more monosyllabic words due to the more flexible syllable structure. Another distinctive property is that Finnish has two and Estonian three different vowel and consonant quantities (lengths) that convey phonemic contrast. Although conversational Finnish and Estonian tend to deviate notably from the formal language, the words of conversational Finnish and Estonian still have more syllables than in English (Table 1).

In case of the Switchboard corpus, we used the original pre-segmented utterance-level signals and manually verified syllable-level transcriptions created by ICSI Berkeley Switchboard Transcription Project. As for the EstPhon and FinDiag corpora, the original dialogue-length recordings were divided into utterance-length signals using the utterance boundary markings available in the corpora transcriptions. Any consecutive utterances having an intervening pause shorter than 500 ms were merged into one utterance. The EstPhon comes with manually verified phone-, syllable-, and word-level transcriptions and contains studio dialogues, lecture monologues and fieldwork dialogues. In the present experiments, we used the entire section of the studio-recorded dialogues. As for the Finnish data, the current version of FinDiag contains two dialogues (four talkers) that had been manually annotated at the syllable level at the time of the present study, and these dialogues were used in the present experiments. All three corpora contain speech from both male and female talkers.

We also wanted to analyze performance on infant-directed speech (IDS) in order to compare against the existing approaches to word

⁶ Note that multiplication (log-summing) across *all* bands would be problematic with wideband speech (0–8000 Hz at 16,000 Hz sampling rate) since many speech sounds contain very little energy in some bands, potentially leading to multiplication by noisy values close to zero.

Table 1

Summary statistics for the four corpora used in the study. Speaker identities (the number of talkers) were not available for the Switchboard ICSI subset but are counted in tens.

	Audio duration (min)	N talkers	N utterances	N syllables	N words	Syllables/s	Words/s	Mean utt. duration (s)
EstPhon	310.77	21	10158	107765	64299	5.60	3.50	1.84
FinDiag	63.97	4	2213	26272	14632	6.27	3.85	1.73
Switchboard	159.48	N/A	3824	46504	37265	4.91	3.95	2.50
Brent	102.63	4	6253	23776	19875	3.77	3.19	0.98
Total or mean	636.85	> 29	22448	204317	136071	5.14	3.62	1.76




Fig. 4. Visualization of signal sonority envelope representations used as the basis for syllabification in the four compared algorithms, shown for a Brent corpus utterance “Let’s snap your shirt”. The amplitudes of each method have been scaled for visualization purposes. Tuning of $f_c = 8$ Hz, $Q = 0.8$ is used for the oscillator model (see Appendix B for explanation).

segmentation using statistical learning and phonetic knowledge. Therefore a subset of the Brent corpus (Brent & Siskind, 2001) that was automatically transcribed by Rytting, Brew, and Fosler-Lussier (2010) was also included in our analyses (this subset was referred to as “Large Brent” in the study of Rytting et al., 2010). Since these forced-aligned transcription data only contain markings on phone segments, we reverse-engineered the word-level timestamps from the phone data using the CMU pronunciation dictionary and the original CHILDES orthography of the utterances. Utterances containing words whose pronunciations were not unambiguously mapped between the orthography and the forced-aligned phones by the pronunciation dictionary were excluded from our analyses. We syllabified the phone transcripts of the Brent corpus with the commonly used rule-based tsylb2-algorithm (Fisher, 1996) that is based on the syllabification rules described in Kahn (1976). Although the quality of resulting linguistic reference is not directly comparable to the manually annotated corpora (due to several sources of uncertainty),⁷ it was sufficiently consistent to provide an overall indication of the algorithm’s performance on IDS and enabled comparison to the study of Rytting et al. (2010), where the forced-aligned data was used as a linguistic reference.

Descriptive statistics for all four corpora are summarized in Table 1. All non-speech segments in all corpora were marked as “empty” and not considered as linguistic units in the evaluation. Although Switchboard contains only narrow-band telephone speech (300–3400 Hz, in contrast to the other wideband-recorded corpora), this limitation was not considered to be an issue since the narrow-band telephone speech still contains sufficient acoustic structure for intelligible speech (due to this, the number of frequency bands in Eq. (2) was $N = 16$ for Switchboard).

Free parameters of all algorithms (see below) were optimized on a separate development set consisting of 500 randomly chosen utterances from each corpus ($4 \times 500 = 2000$ utterances in total) using syllable

segmentation F-score as the criterion (Section 3.3). These development set utterances were excluded from the final test used in the reported experiments. Since the number of optimized parameters in each algorithm was relatively small (max 6), we believe that the size of the development set was large enough to avoid significant overfitting to the data. In addition, for each algorithm, any constant delay between algorithm outputs and reference syllable boundaries was automatically corrected at the corpus level to remove any delays caused by filtering and windowing operations. In practice, this corresponded to a maximum of some tens of milliseconds per algorithm.

3.2. Reference algorithms

We compared the oscillator model against three existing syllabification algorithms: (1) an envelope velocity-based algorithm proposed by Villing et al. (2004), from now on referred to as “VSeg”, (2) a simple amplitude envelope minima detector using ear-like temporal filtering (“EnvMin”), and (3) a sonority-based speech-rhythm estimator developed by Wang and Narayanan (2007) and modified by the current authors for syllable boundary detection purposes (“WN”). Examples of the resulting rhythmic representations in all compared algorithms can be seen in Fig. 4.

3.2.1. VSeg

The basic idea of VSeg is to determine syllable onsets as maximal positive peaks in the velocity (i.e., rate of increase) of a low-pass filtered amplitude envelope of speech. In order to refine the quality of the segmentation, the algorithm ensures that the detected peaks are followed by a potential nucleus, a sonorant sound with prominent energy below 1-kHz (F1-region). In addition, the method inhibits smaller velocity peaks within ± 100 -ms from prominent peaks from being considered as segment boundaries (see Villing et al. (2004), for details). In the original paper VSeg was shown to compare favorably against the classical convex-hull algorithm (Mermelstein, 1975) on read speech and has since been used in, e.g., paralinguistic speech processing (An, Brizan, & Rosenberg, 2013). In the present work, we used a MATLAB implementation of the algorithm using the slight modifications

⁷ For instance, the original lexicon-driven orthography may not reflect phonetic pronunciations, automatic forced-alignment with a speech recognizer may lead to temporal alignment errors, tsylb2 is agnostic to the underlying rhythmic structure and deals only with general phonological rules (see also Marchand, Adsett, & Dampier, 2007).

introduced in the version⁸ included in the AuToBI-toolbox (Rosenberg, 2010).

There are six free hyperparameters in VSeg that were all jointly optimized on the development set using an exhaustive grid search. The tested values included 0.5, 1, and 2 times of the default values in AuToBI and Villing et al. (2004), except for the most important scoring thresholds s_{\min} and s_{\max} , which were searched between 0.3 and 0.9 in 0.1 increments⁹. The resulting optimal parameter values were $s_{\min} = 0.5$, $s_{\max} = 0.6$, $b_{\min} = 0.005$, $b_{\max} = 0.05$, $\nu p_{\min} = 0.005$, and $\nu p_{\max} = 0.2$ (see Villing et al. (2004), for explanation).

3.2.2. Amplitude envelope segmentation

Although very straightforward, direct minima detection from smoothed amplitude envelopes has the potential to yield rhythm-based segmentation of speech without the risk of overfitting the model parameters to a specific language or speaking style. Therefore this approach was also included as a baseline method.

In the current minima-based algorithm (“EnvMin”), full-wave rectification is first performed to the incoming signal, followed by down-sampling to 1000 Hz. The resulting signal is then smoothed in time using a 100-point low-pass FIR filter that approximates the shape of the temporal window of integration in human hearing (see Räsänen & Laine, 2013). In comparison to the standard moving average filter, this minimum-phase filter maintains sharp onset-detection while still providing efficient smoothing of the envelope with approximately 7-Hz cutoff-frequency.

In order to detect syllable boundaries, the envelopes are normalized to a value range of [0, 1] and each local minimum that is preceded by a local maximum of at least δ units higher is marked as a candidate boundary. Boundaries closer than 50-ms to each other are replaced with a single boundary located at the deeper envelope trough of the two candidates. The threshold was optimized to $\delta = 0.10$ using an exhaustive search on the development set with small increments.

3.2.3. WN

Our third reference algorithm is based on a speech-rate estimator developed by Wang and Narayanan (2007). Similar to the present algorithm, WN is based on correlations of the signal energy envelope across multiple frequency bands. The main difference from the oscillator model is that it uses a set of specialized signal processing operations (temporal autocorrelation and filtering) to smooth the signal energy in time instead of using a bank of harmonic oscillators. The algorithm was originally designed for speech rate purposes, estimating the number of syllabic nuclei (sonority envelope maxima) per time unit, rather than detecting potential syllable boundary locations (sonority minima), and therefore the sonority envelope computation was optimized in the original paper to solve the former problem. Still, as the algorithm is a well-tested method for sonority estimation, it was of interest how well it compares against the other syllabification methods if we simply perform minima detection from the resulting sonority envelopes.

Due to the large size of the parameter space and computational complexity of the WN, we used the default settings from Wang and Narayanan (2007) that were originally optimized on Switchboard data with a series of search techniques. In order to detect boundaries from the WN envelopes, a minima detection algorithm equivalent to the EnvMin algorithm was used after the utterance-level envelopes had been scaled between 0 and 1. The detection threshold of WN was optimized to $\delta = 0.1$ using an exhaustive search on the development set with small increments.

⁸ <https://github.com/vsoto/autobi/blob/master/src/edu/cuny/qc/speech/AuToBI/Syllabifier.java>. Accessed on March 4 th, 2015.

⁹ A more detailed search was not computationally tractable in the 6-D space.

3.3. Evaluation

The overall temporal accuracy of the segmentation was measured by checking whether the boundaries of detected acoustic chunks resided within ± 50 ms from an annotated syllable boundary (Obin et al., 2013; Villing, Ward, & Timoney, 2006). This ± 50 -ms window can be seen as a reasonable compromise between accurate boundary detection and the uncertainty from the intra- and cross-annotator consistency at the phone level (see Kvale, 1993) and the uncertainty in defining syllable boundaries during, e.g., fricatives and closures between two syllables. In general, this window length is in similar proportion to the length of the underlying syllabic units as the ± 20 ms window used in phone segmentation is to the average phone lengths (Almpanidis & Kotropoulos, 2008; Esposito & Aversano, 2005; Kvale, 1993; Scharenborg, Ernestus, & Wan, 2007).

Following the standard procedure of speech segmentation literature (see Räsänen, Laine, & Altosaar, 2009), recall was defined in terms of the number of boundaries correctly detected divided by the number of manually annotated boundaries ($R = N_{\text{hit}}/N_{\text{ref}}$), precision as the proportion of hypothesized boundaries also corresponding to true boundaries ($P = N_{\text{hit}}/N_{\text{hypo}}$), and F-score as the harmonic mean of the two ($F = 2 * P * R / (P + R)$). As an additional constraint, each algorithm boundary could only match with the nearest reference boundary in order to prevent double-counting a boundary placed in the middle of a short syllable—a rare but possible event.

Segmentation performance was measured separately for syllable and word boundaries. Since languages such as English contain a large proportion of monosyllabic word tokens in conversational speech (Greenberg, 1999), it was of interest whether rhythmic cues could provide direct access to word-like units without any additional mechanisms such as statistical learning (Saffran et al., 1996). In addition to measuring the accuracy for individual word and syllable boundaries (“boundary performance”), performance of segmenting full units correctly was measured separately by requiring that both boundaries surrounding a manually annotated syllable/word were detected according to the ± 50 ms criterion without any insertions in the middle of the unit (“unit performance”).

Furthermore, in order to reflect the normal infant speech perception task with clear cues for utterance boundaries (cf., Seidl & Johnson, 2006), any missing boundaries at the onsets or offsets were automatically added to score utterance endpoints as correct detections. This is because all compared algorithms are designed to detect local envelope minima that do not necessarily exist at the beginning or ending of an utterance or velocity maxima (VSeg) that do not exist after the last syllable. Since the onsets and offsets are clearly indicated by the envelopes (see Fig. 4) and since temporal ambiguities at utterance boundaries play no linguistic function, we wanted to avoid tuning separate detection mechanisms for these trivial boundaries for each method. On a similar basis, boundaries detected during the non-speech segments preceding and following the spoken utterances were excluded from the analysis (e.g., caused by background noise or crosstalk from another talker). Finally, after initial optimization, all algorithm parameters were kept constant across all corpora in order to get a measure of overall performance across different languages without language-specific tuning.

3.4. Overall segmentation performance

We first report the overall results from our tests, followed by a more detailed analysis of the findings. Statistical significance of the differences between the algorithms or corpora is not reported since all methods are deterministic, but small differences should be reliable since all measures are based on thousands of tokens.

Table 2 shows the results for syllable boundaries for all corpora and compared methods. Tables 3–5 show the corresponding results for full syllables, word boundaries, and full words, respectively. Oscillator

Table 2

Syllable boundary segmentation performance for all four corpora and the four algorithms compared. *P* stands for precision, *R* for recall and *F* for F-score. Highest values across the four methods for all corpora and performances measures are indicated with boldface font.

	Estphon			FinDiag			SWB			Brent			Mean		
	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>
Oscillator	0.71	0.78	0.74	0.77	0.72	0.74	0.74	0.84	0.79	0.63	0.80	0.70	0.71	0.78	0.74
Envelope	0.67	0.71	0.69	0.69	0.64	0.66	0.63	0.72	0.67	0.65	0.74	0.69	0.66	0.70	0.68
Vseg	0.82	0.63	0.71	0.85	0.57	0.68	0.81	0.69	0.75	0.73	0.68	0.70	0.80	0.64	0.71
WN	0.77	0.62	0.69	0.79	0.56	0.65	0.82	0.67	0.74	0.72	0.72	0.72	0.78	0.64	0.70

Table 3

Segmentation performance for full syllables. Highest values across the four methods for all corpora and performance measures are indicated with boldface font.

	Estphon			FinDiag			SWB			Brent			Mean		
	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>
Oscillator	0.51	0.58	0.54	0.51	0.48	0.49	0.56	0.69	0.62	0.41	0.55	0.47	0.50	0.58	0.53
Envelope	0.44	0.48	0.46	0.41	0.38	0.39	0.41	0.51	0.45	0.38	0.45	0.41	0.41	0.46	0.43
Vseg	0.46	0.35	0.40	0.42	0.27	0.33	0.51	0.46	0.48	0.36	0.33	0.35	0.44	0.36	0.39
WN	0.41	0.34	0.37	0.36	0.25	0.30	0.49	0.43	0.46	0.40	0.40	0.40	0.42	0.35	0.38

Table 4

Segmentation performance for word boundaries. Highest values across the four methods for all corpora and performance measures are indicated with boldface font.

	Estphon			FinDiag			SWB			Brent			Mean		
	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>
Oscillator	0.44	0.75	0.55	0.46	0.74	0.57	0.60	0.82	0.69	0.58	0.86	0.70	0.52	0.79	0.63
Envelope	0.42	0.69	0.52	0.42	0.66	0.51	0.52	0.72	0.60	0.61	0.81	0.70	0.49	0.72	0.58
Vseg	0.51	0.61	0.56	0.53	0.60	0.56	0.66	0.69	0.67	0.69	0.74	0.72	0.60	0.66	0.63
WN	0.48	0.61	0.53	0.50	0.59	0.54	0.67	0.67	0.67	0.69	0.79	0.74	0.58	0.66	0.62

Table 5

Segmentation performance for full words. Highest values across the four methods for all corpora and performance measures are indicated with boldface font.

	Estphon			FinDiag			SWB			Brent			Mean		
	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>	<i>P</i>	<i>R</i>	<i>F</i>
Oscillator	0.28	0.54	0.36	0.29	0.49	0.36	0.45	0.69	0.55	0.41	0.67	0.51	0.36	0.60	0.45
Envelope	0.24	0.46	0.32	0.23	0.40	0.29	0.33	0.51	0.40	0.39	0.56	0.46	0.30	0.48	0.37
Vseg	0.25	0.33	0.28	0.25	0.29	0.27	0.40	0.46	0.43	0.40	0.44	0.42	0.32	0.38	0.35
WN	0.23	0.32	0.26	0.22	0.27	0.25	0.39	0.42	0.41	0.44	0.53	0.48	0.32	0.39	0.35

results are reported for $f_0 = 8$ Hz, $Q = 0.8$, and $\delta = 0.01$ that were found to be optimal on the development set (see Appendix B for analysis).

The first main finding is that the oscillator model performs the best across all compared models and is successful at discovering a large number of syllable boundaries in all three languages with $F \geq 0.74$ for the manually annotated data and $F = 0.71$ for the automatically syllabified Brent IDS. However, this means that the model output differs from the reference annotations (i.e., from $F = 1$). This gap shows that sonority-based syllabification, as defined and modeled in the present study, deviates from linguistically defined syllables in several cases and suggests that improved sonority analysis or additional constraints such as lexical or segmental knowledge would be needed to improve the accuracy. Since the oscillator is not especially biased towards over-

segmentation (low precision and high recall) or under-segmentation (vice versa), it seems that the approach is not simply missing some of the boundaries or creating too many of them. Instead, it is simply segmenting the data in a manner different from the linguistic reference, ratifying the potential differences discussed in Appendix A between phonological syllables and acoustic chunks.

As can be expected, the average segmentation accuracy is lower for words than syllables for all algorithms. However, a notable outlier to this pattern is the IDS speech of the Brent corpus where the segmentation performance for words is similar or even higher than the syllabification performance, both in terms of single boundaries and full units. This is also opposite to the pattern in adult-directed conversational speech of the same language (Switchboard data) where syllables are more accurately segmented than words. However, due to varying

annotation procedures, speaker, and acoustic channel characteristics (Section 3.1), the Brent and Switchboard data are not directly comparable, and these results should not be taken as conclusive differences between IDS and ADS. An analogous problem is associated with the comparison of segmentation performance across different corpora. Although it appears that the Finnish and Estonian are somewhat harder to syllabify than English (Switchboard), it is impossible to know whether cross-language differences are caused by (1) actual differences between languages, (2) differences in the annotation principles or annotator consistency, (3) differences in the acoustic properties of the recordings, (4) or due to some other uncontrolled factors involved in collection of conversational speech. Therefore the current results across different languages should be seen as an indication of the existence of useful rhythmic structure in speech, all languages having similarly accessible syllable structure from sonority point of view, but not providing reliable evidence for any detailed between-language comparisons or conclusions.

As for the comparison between the syllabification algorithms, the oscillator model performs the best on average when measured in terms of the average F-score. It appears that VSeg and WN consistently undersegment, having relatively high precision but low recall for the boundaries. Since these algorithms miss many of the boundaries, precision on full units is dramatically lower as the hypothesized segments (two subsequent boundaries) rarely align with both onset and offset of a syllable or word. This may be caused by the large amount of temporal smoothing performed during the sonority estimation, causing loss of information that seems to be relevant for the detection of some of the subtler syllable boundaries in conversational speech. However, it should be noted that the high precision and average recall of the WN can be a desired characteristic in speech rate estimation (the original task of the algorithm) as long as this trade-off leads to more consistent syllable counts within a time window of interest. The goal in speech rate estimation is to have a reliable correlate of the speaking rate, and therefore any systematic error in the nuclei count can be easily compensated for if absolute syllable counts are required (see Wang & Narayanan, 2007).

We created random baselines for the full test set by using the same number of hypothesized boundaries as were produced by the oscillator algorithm, but uniformly randomizing the temporal positions of the boundaries in the utterance. This baseline produced an average F-score of 0.52 (± 0.03 SD across corpora) for syllable boundaries, 0.24 (± 0.03) for syllables, 0.46 (± 0.04) for word boundaries, and 0.22 (± 0.03) for words, confirming that all compared algorithms were performing substantially better than a random segmenter with an informed estimate for the number of syllables in an utterance.

3.5. Distribution of errors

In order to get more detailed picture of the syllabification process, Fig. 5 shows the distributions for the temporal difference between the hypothesized and reference boundaries in all four corpora. On average, 53% of the boundaries are within ± 20 ms, 72% within ± 50 ms, and 88% within ± 100 ms from the references, while the corresponding numbers for randomized baseline using the informed number of sonority maxima per utterance are 22%, 46%, and 69%. There are also two other observations to be made. First, the distributions are asymmetric. If a hypothesized boundary is within ± 50 -ms from a reference boundary, it is ruled out from further matches with other reference boundaries. In this case, if the next reference boundary is missed, the nearest hypothesized boundary is likely located *after* the reference due to the exclusion of the preceding one. Second, every language save the automatically-syllabified Brent data has a small peak in the distribution around 40–50 ms before the reference boundary. Closer analysis of the EstPhon data (Appendix C) revealed that this peak is caused by systematic mismatch between manually and automatically syllabified boundaries in cases of stop sounds (e.g., [p]). Reference boundaries

were always placed at the onset of the stop burst, whereas the sonority minimum could be anywhere during the closure preceding the stop.

We also analyzed the parametric sensitivity of the oscillator model (Appendix B) and its segmentation performance as a function of phones flanking the syllable boundary (Appendix C) and as a function of the part-of-speech of the words (Appendix D). Full details are available in the appendices but we summarize the conclusion of these analyses below. The sensitivity analysis shows that the segmentation performance is not highly dependent on exact tuning of the oscillator characteristics, with a wide range of tunings leading to reasonable accuracies. In addition, language-specific optimization of the parameters gives only a minor boost from the average syllable boundary F-score of 0.74 (Table 2) to $F = 0.76$, supporting the idea of sonority as a universal (but incomplete) cue to syllabic structure in speech. The error analyses show that the typical syllabification errors occur at syllable boundaries where neither the offset of the preceding syllable or the onset of the following syllable contains a low-sonority phone, aligning with the findings that humans also tend to make syllabification errors at boundaries flanked by segments with similar sonority ratings (Price, 1980, and references therein). This analysis confirms that the acoustically grounded sonority sequencing principle (SSP) works well for languages with universally prevalent CV-like syllable structures, especially when the consonant is a stop or a fricative sound (see also Fig. A1 in Appendix A). Finally, the part-of-speech analysis shows that syllables of English IDS content words such as nouns and verbs tend to be segmented with above-average accuracy while syllables in function words are much more difficult to extract based on sonority cues alone, suggesting that the content words have more salient rhythmic structure.

3.6. Conclusions from the experiments

The experiments show that the sonority envelope has a similar relationship to the underlying syllabic structure in English, Finnish and Estonian, supporting the idea that SSP as a cross-linguistic property of speech is also reflected in acoustically grounded measures of sonority. Moreover, all methods using sonority-like measures for syllabification were relatively successful in all languages, with the oscillator model leading to the best overall performance. This pattern of results suggests that the exact mechanisms of sonority tracking are not critical for the extraction of useful rhythmic information from speech, but that the best results are obtained when the analysis is carried out using a system mimicking the basic characteristics of auditory perception for amplitude modulation on different frequency bands.

The acoustic chunks resulting from the oscillator model are still far from a perfect linguistic syllabification, however. Many of these differences are systematic (see Appendices A and B), largely because the reference syllabification always aligns with phone segment boundaries while the actual sonority minimum can occur in the middle of phonological segments. In addition, some syllable boundaries, typically flanked on both sides by highly sonorous sounds, are simply too subtle to be detected by the present approach. These failures leave open the question of whether the current models are too simple to account for fully human-like perception of syllabic rhythm in speech. Another possibility, however, is that phonology-driven syllabic annotations reflect linguistic structure that diverges from perceptual processing of speech and depends on formal language training (see Port & Leary, 2005).

In sum, despite the discrepancies, we find that syllable-like acoustic chunks can be extracted in a non-language-specific way, suggesting that these elements could be used in early language acquisition. The differences between theoretical and perceptual syllables may be resolved using information from other linguistic levels, or they may persist in interesting and testable ways (see Section 4.4).

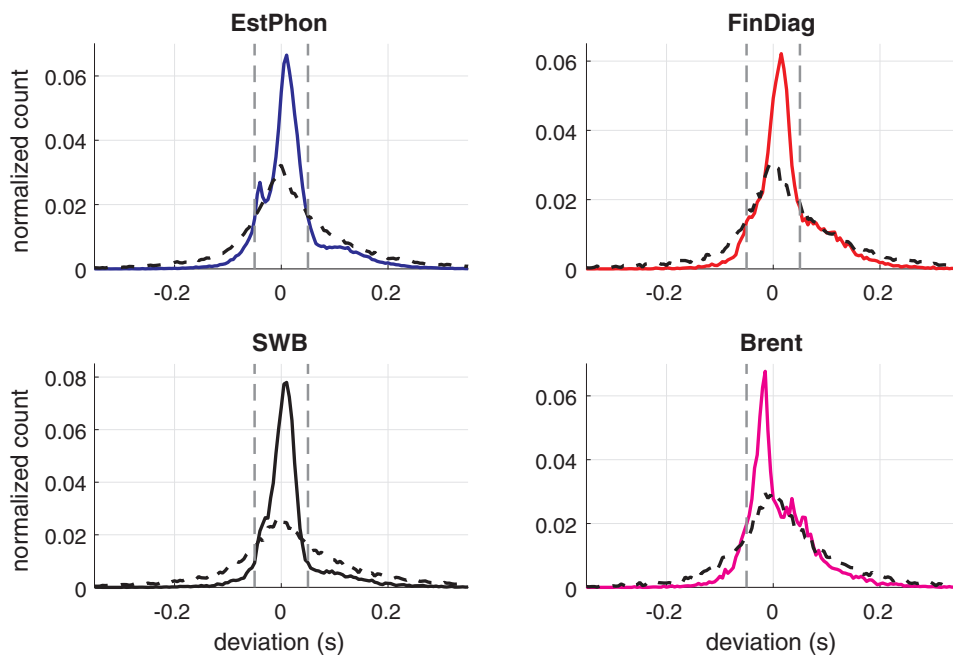


Fig. 5. Distributions for the deviation between hypothesized and reference syllable boundaries ($t_{\text{hypo}} - t_{\text{ref}}$) for the four corpora using the oscillator model. Boundaries from random baseline are shown with black dashed lines. The allowed deviation of ± 50 ms for a boundary to be considered as correctly segmented is indicated with vertical grey lines.

4. Discussion

The concept of a syllable is central to many theories and models of language acquisition. Syllables are often considered as units of speech that are accessible when little, if any, phonological learning has taken place and before word learning begins, making them highly relevant for early language acquisition. Previous work has been vague, however, regarding both how infants could access the syllabic structure of speech and how reliable and accurate their syllabification is. This ambiguity extends even to phonetic and phonological theories, which have struggled to find generalizable characterizations for these seemingly simple and salient units of speech (see Krakow (1999), for a review). By describing a strategy for extracting syllable-like acoustic chunks from raw acoustic input, our work ties the notion of the syllable to the acoustic structure of speech.

At a low level, a perceptual basis for syllables would not only explain why syllables are perceptually salient (by definition), but also: why access to syllabic structure does not require formal language training (e.g., Morais et al., 1986), why syllabic structure shows up already as early as the primary auditory cortex (Ahissar et al., 2001; Giraud & Poeppel, 2012), and why talkers (who are not professional linguists) tend to utilize overt or covert articulation to obtain an impression of the constituent syllabic structure when asked to syllabify something.¹⁰ In fact, perceptual primacy of syllables is implicitly assumed in infant language studies where syllables are considered to be accessible before the actual language learning starts. This perceptual basis for syllables may evolve during language development, with child's representation gradually evolving from initial acoustic-based syllable-like chunks into more formal linguistic entities that satisfy emerging constraints from multiple levels of representation.

In the present study, we defined a proxy for a pre-linguistic proto-syllable as a chunk of speech determined by signal sonority fluctuation over time, accessible using generic auditory mechanisms. We created a computational model for delineating such chunks from the acoustic signal, unconstrained by the underlying linguistic structure of the language, and investigated how the outputs from this model and from a

number of previously proposed algorithms are related to the underlying phonological syllables. As a result, we were able to study how well such chunks, identifiable from conversational speech before any language-specific learning takes place, correspond to syllables and words of the language.

Experiments in three different languages indicated that the sonority envelope of speech was highly correlated with phonetically-annotated syllabic structure across different language families and syllable distributions, supporting the idea of the sonority sequencing principle as a language-universal property of speech and the idea of syllables as accessible (if imperfect) units in the early stages of language learning. More detailed analysis revealed that there were systematic discrepancies between the acoustically-determined proto-syllables and the standard theoretical syllable boundaries. While the perceptual proto-syllables are useful as a starting point, the standard conception of syllables arises from the combination of perceptual information and other linguistic structure (cf., Cummins, 2012).

Our results also provide indications that infant-directed speech (IDS) may be more robust to the exact tuning of the rhythmic analyzer in contrast to adult-directed speech (ADS). Unfortunately, proper verification of this finding is not possible without a dataset of IDS and ADS with matched linguistic content and with manually verified annotations at the syllabic level—a resource that, to the best of our knowledge, is not currently available. The analysis with POS categories (Appendix D) also revealed that content words in English IDS are especially accessible with rhythmic cues whereas only a minority of function words are segmented as isolated units. This aligns well with the finding that children tend to acquire content words earlier (Fenson et al., 1993; Frank, Braginsky, Yurovsky, & Marchman, 2017), and supports the idea that concrete nouns and verbs may be easier to learn not only due to their significant referential support (Räsänen & Rasilo, 2015), but also because they are rhythmically more prominent.

4.1. Relation to other segmentation studies

To our knowledge, this is the first study to address in detail and cross-linguistically the acoustic and perceptual characteristics of acoustic chunks that are potentially accessible to a pre-linguistic learner from realistic conversational speech, previous studies focusing on read speech or individual languages (e.g., Hyafil et al., 2015; Mermelstein, 1975; Obin et al., 2013; Villing et al., 2006; see also Leong and

¹⁰ For experienced language users, learned constraints from speech production are also likely to contribute to a more detailed assignment of consonantal units to the syllables. Note that this type of knowledge is only emerging for pre-linguistic infants.

Goswami (2015), for related work). The overall syllabification results obtained for different languages are in line with the earlier studies on English read speech that have reported F-scores in the range of 0.7–0.82 (e.g., Mermelstein, 1975; Obin et al., 2013; Villing et al., 2006).

As a further reference, Rytting et al. (2010) applied the word segmentation algorithm of Christiansen, Allen, and Seidenberg (1998) to the Brent data, testing both citation-form phones and posterior phone probabilities from a supervised speech recognizer as an input to the algorithm. They obtained $F = 0.66$ for word boundaries and $F = 0.29$ for words using the canonical phones and $F = 0.58$ for word boundaries and $F = 0.18$ for words when using the probabilistic input that is more affected by the acoustic variation in speech (but still uses supervised training to significantly reduce non-phonological variation in the resulting representation). Although the present results are not directly comparable due to the differences in evaluation,¹¹ the oscillator model achieves significantly higher word segmentation performance ($F = 0.70$ for word boundaries and $F = 0.51$ for full words using the default language-generic parameters for syllabification and $F = 0.77$ and $F = 0.51$ with parameters optimized for word segmentation; not shown separately) on the same data while facing the full complexity of the acoustic speech and without using any linguistic knowledge to constrain the analysis. Note that word segmentation performance by the oscillator model is still far from the results obtained using phoneme- and canonical syllable-level representations in state-of-the-art statistical algorithms (e.g., Phillips & Pearl, 2012). However, these models assume that the speech is already correctly syllabified and categorically perceived, and they therefore have a fundamentally different starting point for the discovery of language structure in the input.

4.2. Relationship to speech perception research

How does an acoustic chunk representation fit into the larger problem of speech perception? Our findings support the basic idea that entrainment to speech rhythm can divide the speech input into frames that have a consistent relationship to the underlying linguistic content. But a syllabification process like the one we propose is likely to provide probabilistic rather than categorical cues. As seen in the oscillator model's performance, the rhythmical, phonological, and lexical levels do not always align perfectly. In fact, it is unclear how any system operating in a purely bottom-up manner could determine fully accurate phonological syllable boundaries due to these conflicts (see Fig. A2 in Appendix A for a concrete example). As long as the brain is assumed to be sampling the speech in signal-driven rhythm, the presence of alignment uncertainty suggests that the rhythmic representations should be treated as probabilistic, not as hard decisions into indivisible syllabic segments, before further processing. How this type of probabilistic interpretation fits to the neurophysiological level is currently unclear, as these models tend to assume that the front-end syllabic chunking is driving the rest of the processing (e.g., Giraud & Poeppel, 2012).

A probabilistic interpretation of syllabic structure is also more flexible in the face of ambiguities involved in real speech, and the range of information types humans use to understand speech. Any categorical decisions on the hypothesized signal structure should not be made based on individual cues, but by considering the entire speech decoding problem as a whole, and integrating many, often conflicting, cues. For instance, hypothesized intentions of the speaker (e.g., Frank, Goodman, & Tenenbaum, 2009), referential context (e.g., Räsänen & Rasilo, 2015, and references therein), phonotactic regularities (e.g., Adriaans &

Kager, 2010), rhythm of the syllables (Cutler & Norris, 1988), and spectral and statistical cues to phone-like segments (e.g., Räsänen, 2014; Scharenborg et al., 2007) and words (e.g., Räsänen, 2011; Saffran et al., 1996) are all potentially informative of the structure and content of the message.

In general, the use of any individual cue should be based on the reliability of that cue, and the final interpretation should depend on the demands of the ongoing perceptual task. The set of available cues and their reliability changes constantly over the course of language learning. Features that are not language-specific, such as general rhythmic cues, may be central to early speech parsing before losing ground to later-emerging but more reliable cues, including language-specific ones (Bertoncini & Mehler, 1981; Bertoncini et al., 1988; Bijeljac-Babic et al., 1993; Eimas, 1999; Jusczyk & Derrah, 1987; Jusczyk et al., 1990; Jusczyk et al., 1995).

Finally, our results show that a syllable-based language learning bootstrapping strategy can be useful not only for syllable-timed languages such as Finnish or French, but also for stress-timed languages such as English. Earlier studies have found out that syllabic rhythm is used in word recognition by French but not by English talkers (Cutler, Mehler, Norris, & Segui, 1986, 1992) and that English speaking infants start to consider trochaic (strong-weak) stress patterns as holistic units already 7.5 months of age (e.g., Jusczyk, Houston, & Newsome, 1999). In contrast, French infants are capable of segmenting speech into syllables already at the age of 6 months (Nishibayashi, Goyet, & Nazzi, 2015) and keep doing so at least until 12 months of age (Nazzi, Iakimova, Bertoncini, Frédonie, & Alcantara, 2006). This body of evidence led Nazzi et al. (2006) to put forward the hypothesis that early language parsing strategies may depend on the rhythmic class of the language in question, leading to crosslinguistic differences in the use of prosodic cues for segmentation. The present computational analysis shows that rhythmic segmentation can directly provide a large number of correctly segmented word forms also for English (> 50% in Brent IDS) and no major differences in syllable segmentation performance were observed between languages from different rhythmic families. While this analysis suggests that rhythmic syllable-like structure is, in principle, also useful in English, the present findings do not exclude the possibility that learners in stress-timed languages may also switch from individual syllables to trochaic stress patterns early in the development. The high prevalence of monosyllabic words or trochaic multisyllabic words in English might simply mean that there is no pressure for the emergence of a separate and intermediate syllabic decoding process, as direct lexical access from the contents of a trochaic pattern is usually a feasible strategy.

In the future work, the model could be extended to also include other generic auditory cues such as sound onsets (spectral change) that the auditory system is known to be sensitive of. In addition, other cues to linguistic structure, such as statistical regularities at various levels, could be gradually learned by the model with the help from these language-independent auditory segmentation mechanisms. None of these additional capabilities would have to be developed from scratch, as there is a large number of existing models and algorithms for spectral change – based segmentation and statistical learning at various levels (see Section 1). In general, it would be useful to have a model that could replicate the developmental shifts in speech segmentation strategies observed in different languages (see also Thiessen & Saffran, 2003) while operating on raw acoustic input with minimal a priori assumptions regarding properties of different languages.

4.3. Validity and generality of the present findings

The present study suggests that the picture of pre-linguistic syllables as units that always perfectly align with phonological and lexical units might be oversimplified. It remains unclear whether this disparity between the representations is the result of an oversimplified model that uses a limited part of the perceptual information, an aspect of early

¹¹ The model of Christiansen et al. (1998) has access to phone-level signal frames and therefore word segmentation performance was evaluated in Rytting et al. (2010) by determining whether a boundary was between two phones belonging to two neighboring words. In contrast, all present algorithms produce boundary hypotheses in continuous time. Since phone boundaries were not available for all corpora, we used the ± 50 ms accuracy criterion used in the previous syllabification literature (see Section 3.3).

segmentation that disappears as the learner incorporates information from different linguistic levels and structures, or a distinction that persists even for adults.

As with any simplified computational model, it is impossible to determine the actual rhythmic segmentation performance achievable by humans. Instead, the present experiments reveal a *lower bound* for performance that is achievable with a simple perceptual mechanism for tracking sonority variation in speech. Actual human performance may be higher due to experience-based optimization of the perceptual system with respect to language-specific rhythmic patterns (cf., Mehler et al., 1996), even prenatally (e.g., Bortfeld, Shaw, & Depowski, 2013), or simply more efficient or intelligent cue processing by the human auditory system. The current investigation has focused solely on the sonority envelope in speech in isolation from other segmental and suprasegmental properties of speech. Cues such as fundamental frequency contours, spectral tilt, and spectral features are also generally concurrently available for perceptual processing of speech. Even though the current study has not investigated any language or talker-specific adaptation, it is likely that such learning is already involved in early rhythmic perception of speech. This means that the actual infant rhythmic segmentation performance in terms of syllables and words could be, in principle, much higher than the one described here.

Another issue to consider is the method of evaluating the model. The appropriate evaluation depends on whether human syllables are primarily perceptual units or whether they are constrained by an underlying symbolic linguistic structure. If syllables are phonologically constrained, then an ideal syllabification process should align accurately with manual annotation, only limited by the inter-annotator agreement rate. On the other hand, if we assume that syllables are primarily perceptual units, then perhaps another evaluation standard is needed. This question is also further complicated by the potential issue arising from the formal language instruction and phonetic training of the annotators that may bias the annotation process away from what human perception is actually doing (see, e.g., Port & Leary, 2005; Ziegler & Goswami, 2005). Even without such (somewhat controversial) considerations, objective evaluation of the present model against real “auditory” syllables is not possible without an oracle auditory reference of what is actually represented in the brain of the listener, and knowledge of how these representations might interact with other aspects of language. For now, the most informative way to analyze and communicate the findings is to use standard linguistic conventions. However, in the end, a proper test for the modeling work should be based on behavioral experiments.

As mentioned earlier, another major limitation of the present study is related to the comparability of findings across different corpora, which have both meaningful (language, IDS/ADS) and incidental (recording quality, annotations) differences for the language learning problem. All four corpora represent speech from natural communicative situations, but several factors may contribute to the observed results. First of all, the overall sound quality of the utterances varies greatly from recordings in uncontrolled environments and background noises (Brent) to narrow-band telephone conversations (Switchboard) and studio-recorded dialogues (EstPhon, FinDiag). In addition, Brent, EstPhon and FinDiag contain relatively few talkers that necessarily represent only a minor subset of all dialects in each language, limiting generalizations that can be made from such populations. Linguistic content and talker-dependent factors are not normalized between English ADS and IDS, making an accurate comparison of syllabic characteristics in the two speaking styles impossible. Syllable annotation strategies in different corpora also differ, with some manually annotated from the start (FinDiag, EstPhon), while others are automatically syllabified, either with (Switchboard) or without (Brent) manual correction. All these factors limit the existing cross-corpus analyses to a qualitative level. Well-controlled speech data with consistent annotations would be needed to analyze more detailed differences between the languages and speaking styles.

Finally, a technical limitation in the present model is the fixed set of oscillator parameters across all speech input. Although this could be seen as an inherent property of the “perceptual system” similarly to the assumedly fixed modulation sensitivity in the human auditory system (e.g., Dau et al., 1997), the model cannot account for truly oscillatory behavior in speech if this regularity is only observed locally (e.g., utterance level) instead of the global level. As the model has to account for all speech encountered at different speaking rates, strong of constraints with respect to rhythmic regularity (small bandwidth) can lead to a situation where the oscillator cannot react to variation in speaking styles and syllable lengths. As said, this does not exclude the possibility that the lengths of the syllabic units would be somehow constrained by the preceding speech content. However, more complicated adaptive oscillators would be needed to understand whether such local regularities actually occur in speech rhythm, and whether they can be used to constrain and improve the segmentation process even further. In this work, we preferred model simplicity to technical ambiguity, yet still obtained language-independent state-of-the-art performance in automatic syllabification of speech.

4.4. Model predictions

Even though infants seem to learn cross-syllabic statistical regularities independently of whether the resulting chunks are real words of the language (e.g., Ngon et al., 2013), earlier studies also show that statistical regularities crossing intonational boundaries are not captured by learners (Shukla, Nespor, & Mehler, 2007; Shukla, White, & Aslin, 2011) and that stress patterns also affect learning (Hay & Saffran, 2012). These findings suggest that statistical learning is at least partially constrained by the suprasegmental structure of the input instead of operating independently of prosodic parsing. However, it is unclear whether rhythmic structure can also constrain learning, like intonation contours can. If learners are using a sonority envelope-based chunking strategy as suggested by the neurophysiological models (Ghitza, 2011; Giraud & Poeppel, 2012), this strategy would give them access to many correct syllable- and word-like units and explain why they perform so well in standard artificial language learning experiments (cf., Fig. A1 in Appendix A) at a very young age. However, in this case, perceptual chunking into disjoint syllable-like units would also handicap the processing for words that are consistently chunked erroneously, or vary across contexts, in comparison to words with coherent rhythmic forms. For instance, short words or inflections that tend to merge rhythmically to the neighboring words would be expected to be less familiar to infants than rhythmically more distinct syllables or words when both types are equated in terms of their frequency of occurrence. In contrast, words with a clear CVCV-like structure would be expected to be more familiar to the learners due to their rhythmic distinctiveness, especially if they also tend to occur in contexts with clear sonority contrasts with the flanking speech sounds.

A prediction from the rhythmic chunking viewpoint is that the frequency of occurrence and the consistency of rhythmic segmentation across typical contexts would both contribute to the statistical learnability of different phone(me) sequences in speech. This could be tested by applying the model to a sufficiently large corpus of infant-directed speech, extracting a set of frequency-matched phoneme strings with a range of different segmentation consistencies, and testing infants’ listening preference for such tokens (c.f., Ngon et al. (2013), who used a similar approach for testing words against non-words with the same statistical characteristics). More detailed experimental manipulations could be then used to probe which phoneme grouping mechanism best explain the variability in the infant preferences, i.e., whether sonority-based segmentation explains more of the variance than phonological syllabification when they do not align, or whether infants’ listening preferences are totally independent of the syllabic frames and only depend on the recurring strings of phonemic segments. Another test would be to compare intra-syllabic and cross-syllabic phonemic strings

in an artificial language where high-probability phonemic strings would be either aligned or misaligned with the rhythmic (syllabic) units. If rhythmic constraints are important for perception, rhythmically aligned phonemic content should lead to superior learning even when controlling for phonotactic regularity and frequency of occurrence in the stimuli.

5. Conclusions

The present study suggests that pre-linguistic access to syllable-like structures – acoustic chunks – in conversational speech could provide a useful starting point for lexical learning using general auditory mechanisms. We have shown that the rhythmic syllabification works in languages with different stress patterns and different syllabic structure, including in English, whose syllabic structure is considered more complicated than many other languages'. Small but consistent differences exist between perceptual and ideal phonological syllable structure within our model, which may either be resolved through information at other linguistic levels or represent a real difference between the phonological concept of the syllable and its mental representation. Deviations are the smallest for the typical CV-CV-type stimuli used in many artificial language learning experiments (e.g., Saffran et al., 1996), supporting the implicit assumption of these studies that this type of syllabic organization is accessible to a pre-linguistic infant, but also warning that many other phonological syllable structures are less prominent in natural speech when modeled in terms of sonority variation.

Appendix A. Some limits on acoustic rhythmic segmentation

Segmentation based on physically grounded measures of sonority is necessarily ambiguous in certain contexts. As Jones (1960) already noted, it is often impossible to define the limits of a syllable due to the lack of an exact point of minimum prominence between two syllables. Neighboring words, especially short function words, may merge together or undergo substantial reduction at faster speaking rates (see Greenberg (1999), for an analysis). Normal speech also contains words and collocations that exhibit ambisyllabicity, where phonemes may reasonably be claimed to belong to both of two neighboring syllables. This means that acoustic cues to syllabic structure do not necessarily align with the lexical boundaries in speech, a concern already raised by Mehler (1981) in the context of infant perception. It has also been questioned whether syllables (or any other units of language) should be considered to have clear boundaries at all instead of being a succession of rhythmic peaks (e.g. Leoni, 2015). Ambiguity is not usually a problem in speech consisting of simple consonant-vowel (CV-CV) alternation at a regular speaking rate as exemplified by Fig. A1 (cf., the stimuli used in the majority of statistical learning experiments such as Saffran et al., 1996). However, the situation gets more complicated in normal speech (Fig. A2) where such isochrony does not generally exist (e.g., Nespor, Shukla, & Mehler, 2011; Wagner, 2008), and where coarticulatory patterns can also extend across word boundaries.

In the existing language acquisition literature, there has been little discussion on the potential impact of the complexity of the underlying syllabic




Fig. A1. An example of sonority-based rhythmic structure for stimuli in a typical artificial language experiment that uses CV-CV-CV word structure synthesized with MBROLA synthesizer (Dutoit, Pagel, Pierret, Bataille, & van der Vreken, 1996). The top panel shows the original signal waveform while the bottom panel shows a sonority envelope derived from the signal using the method described in the present paper. The red and black vertical lines show the original and automatically detected syllable boundaries, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Based on the success of our model, more work is needed in order to understand whether young children do in fact use sonority-based rhythmic units as indivisible perceptual chunks, as probabilistic cues to linguistic structure, or simply as carriers of information that are flexibly analyzed in terms of sub-syllabic constituents already from the early stages of the language acquisition process.

6. Supplementary material

A MATLAB implementation of the oscillator-based syllabification algorithm is available for download at <https://github.com/orasanen/thetaOscillator>.

Acknowledgements

The authors would like to thank Pärtel Lippus for providing the Estonian Phonetic Corpus of Spontaneous Speech, Mieta Lennes for the FinDialogue corpus, and John Pate for providing the forced-aligned phone transcriptions for the Brent corpus. In addition, the authors are grateful to Martti Vainio, Juraj Šimko, Antti Suni, and Unto K. Laine for useful discussions and comments related to this work. O. Räsänen was funded by the Academy of Finland project “Computational modeling of language acquisition” (no. 274479). G. Doyle and M. C. Frank acknowledge NSF #1456077 and M. C. Frank acknowledges NSF #1528526 for support. A MATLAB implementation of the oscillator-based syllabification algorithm is available at <https://github.com/orasanen/thetaOscillator>.

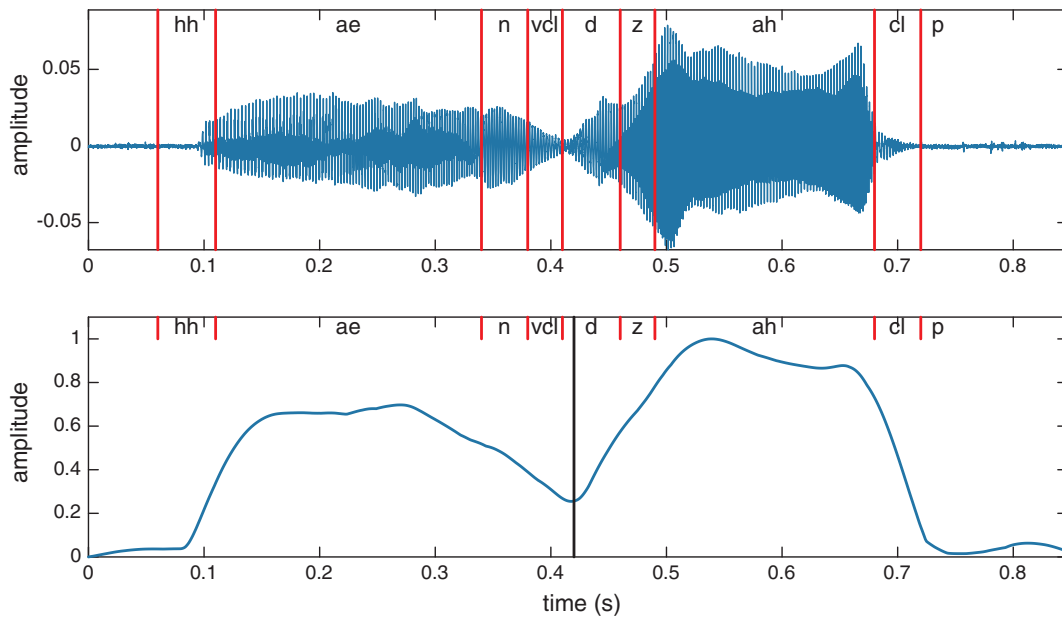


Fig. A2. A more realistic example of sonority-based representation in conversational speech. IDS utterance “hands up” from the Brent corpus (Brent & Siskind, 2001) with the original waveform (top) and the corresponding sonority curve (bottom). Corresponding phone segments are shown with red vertical lines (as automatically annotated by Rytting et al., 2010). Note that [vcl] and [cl] stand for voiced and unvoiced closures of the following stop bursts [d] and [p], respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

or surface acoustic structure on the perceptual parsing of the speech input into some representational chunks. As the above examples illustrate, transparency of the underlying syllabic structure can vary in the surface properties of speech depending on the complexity and style of the input, and this may bear consequences to the design and interpretation of some behavioral studies and models of infant speech perception and language learning.

Appendix B. Parametric sensitivity of the oscillator model

In the experiments in this paper, the qualitative behavior of the oscillator was characterized (and parameterized) using two standard parameters: the center frequency f_c of the oscillator and the so-called Q-factor of the oscillator that depends on f_c and the bandwidth Δf :

$$Q = \frac{f_c}{\Delta f} \tag{B1}$$

When $Q < 0.5$, the oscillator is called over-damped and simply performs temporal smoothing of the variation in the amplitude envelope used to drive the oscillator. In contrast, $Q > 0.5$ is an under-damped oscillator that, when excited, exhibits oscillatory behavior around the center frequency with a decaying amplitude, a larger Q meaning slower decay. The special case of $Q = 0.5$ is called critical damping. In this case, the oscillator tracks the input envelope modulations at rate f_c as closely as possible without causing endogenous oscillatory behavior ($Q > 0.5$) or extra smoothing over time ($Q < 0.5$). Note that the damped oscillator is a true non-decaying oscillator only with $Q = \infty$ ($\Delta f = 0$). With any other Q-factor, it resonates with excitation and decays gradually to zero after the excitation ceases.

In our experiments (see Section 3 for the main results), different values of Q and f_c were tested in order to understand how tuning of the model




Fig. B1. Segmentation performance (the mean F-score across all four corpora) for full syllables (top) and syllable boundaries (bottom) measured as a function of the oscillator model parameters.

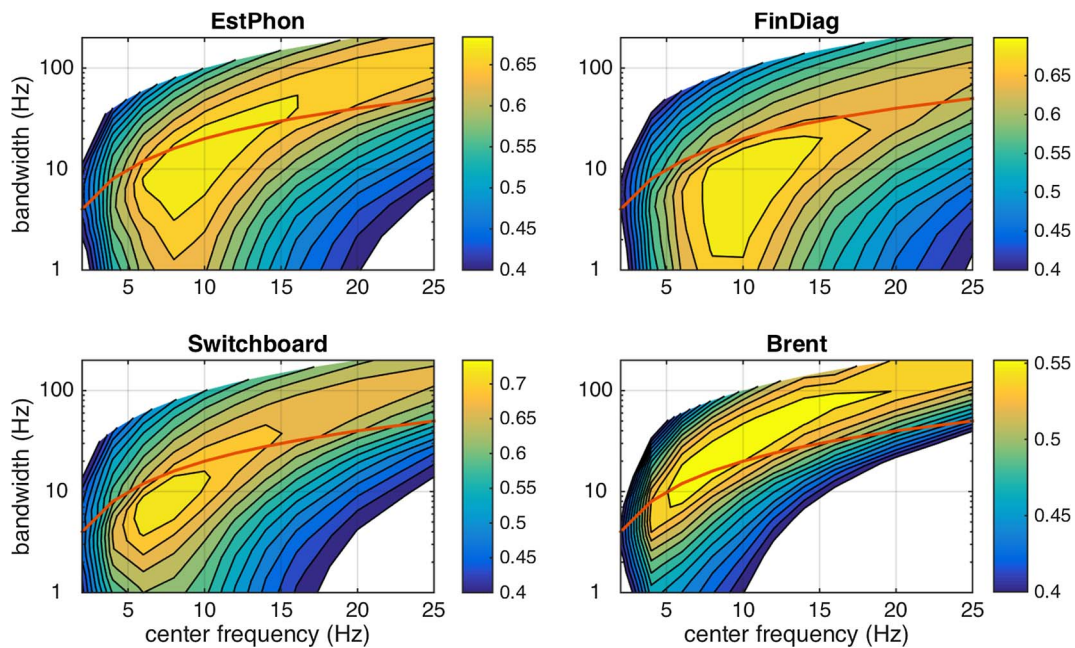


Fig. B2. Corpus-specific performances (F-scores) for utterance-internal syllable boundaries as a function of the center frequency and bandwidth of the oscillator. Each color change from yellow towards blue corresponds to a 5% relative decrement in F-score from the maximum performance when the best performing detection threshold of $\delta = 0.01$ is used. Critical damping of $Q = 0.5$ is shown with the red curve. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

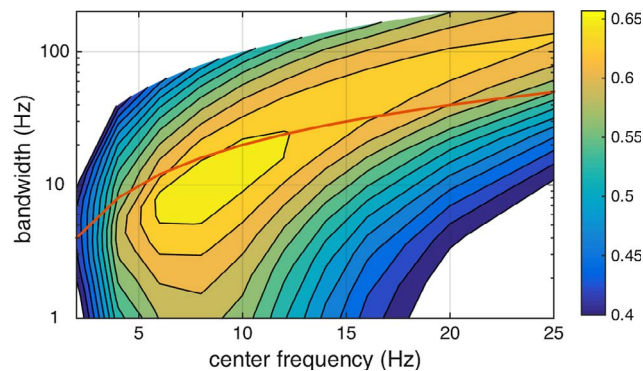


Fig. B3. Performance (F-score) for utterance-internal syllable boundaries as a function of the center frequency and bandwidth of the oscillator, combined across all four corpora. Each color change from yellow towards blue corresponds to a 5% relative decrement in F-score from the maximum performance when the best performing detection threshold of $\delta = 0.01$ is used. Critical damping of $Q = 0.5$ is shown with the red curve. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

impacts the overall performance and whether there are notable language-specific differences in the optimal parameters. A grid evaluation across all parameters (oscillator center frequency, oscillator bandwidth, and detection threshold) was performed using all the data from the four corpora. Fig. B1 shows the overall F-value for syllables (top) and syllable boundaries (bottom) as a function of these parameters when using three different detection thresholds of $\delta = 0.001$ (more segments), 0.01 (medium), and 0.1 (less segments). In order to further illustrate how important exact tuning of the oscillator is in each tested language, Fig. B2 shows the performance for utterance-internal syllable boundaries (i.e., without the trivial onsets and offsets) as a function of the parameters when Q-values are converted to oscillator bandwidths. Fig. B3 shows the same visualization for F-scores averaged across all corpora.

There are three main findings from the analysis. The first finding is that the overall performance is not critically dependent on the exact tuning of the oscillator as long as it covers the relevant frequency range for typical syllable rates (approx. 8 ± 5 Hz). This also corresponds to the range of amplitude modulation frequencies that the human auditory system is sensitive to (e.g., Dau et al. (1997), and references therein) and supports the idea that syllabic rhythm should be perceptually accessible to a pre-linguistic infant.

Second, looking at the smallest bandwidths that are still performing within 5% from the optimal F-score (bright yellow areas in Fig. B2), a relatively narrow tuning still works well for the syllable-timed Finnish (approx. 8 ± 1.6 Hz), followed by the mixed-timed Estonian (approx. 8 ± 3.2 Hz), whereas the stress-timed English requires the broadest tuning with a lower center frequency (approx. 6 ± 4 Hz for Switchboard, 6 ± 8 Hz for Brent). This reflects the complex syllable structure of English, containing a range of syllable lengths and rates in that all should be covered by the segmentation mechanism. However, the present approach cannot exclude the possibility that even English still contains temporally local rhythmic regularities where narrow (adaptive) tuning could function well. What the current results show is that narrow tuning of the perceptual system for specific syllable-rates is more feasible in Estonian and Finnish than in English, but also that such tuning is not required for good

syllabification performance even in the former two languages, as revealed by good performance with larger bandwidths.

The third main finding is that the IDS speech in the Brent corpus has a significantly larger tolerance for variation in the system parameters for syllabification in comparison to the ADS in Switchboard, preferring broadly tuned oscillators that simply act as low-pass filters for the instantaneous sonority values. On the other hand, the overall performance for syllable-internal boundaries is notably lower in the IDS data (Fig. B2, bottom right) than the overall performance with the trivial utterance onsets and offsets also included in the analysis (Table 2), suggesting that much of the IDS performance is due to the shorter utterances and therefore a higher proportion of trivial boundaries. However, due to the reasons discussed in Section 3.4, any systematic differences between ADS and IDS cannot be addressed with the current data sets.

Finally, if the oscillator parameters are optimized to language-specific performance, the mean F-score across all languages is 0.76 for syllable boundaries and 0.60 for syllables. This is only slightly higher than the performance when using the fixed set of optimal parameters across all languages (Table 2). There is also a range of oscillator characteristics in each language for which similar performance is achieved, showing that the present type of syllabification is a relatively robust process and invariant to language specific rhythmic details, including the possible differences in the studied stress- and syllable-timed languages.

Appendix C. Analysis of phones at segment boundaries

In order to better understand the errors in the segmentation, we analyzed the recall of phone boundaries as a function of the phones neighboring the gold-standard syllable boundaries. The analysis was only carried out on EstPhon and Brent data since manually verified phone-level annotation was not available for the other two corpora. Instead of looking at aggregate measures across sound classes, we opted to analyze accuracy as a function of individual phones due to the large number of samples available. All phones pairs spanning a syllable boundary that occurred more than 50 times in the Brent data or 100 times in the EstPhon data were included in the analysis. Table C1 shows the phone-specific recalls for EstPhon and Table C2 for the Brent corpus.

In the EstPhon phone notation, x_v refers to phone x with voicing, x' to x with palatalization, x_+ is x with advanced tongue root, x_o is a lowered x vowel, and x_k is x with creaky voicing. Note also that EstPhon phones are manually annotated whereas Brent phones are from the automatic forced-alignment from Rytting et al. (2010) after they have been converted to a more limited set of sounds accepted by the tsylb2 algorithm that we used for syllabification of the corpus. Since tsylb2 always groups stop closure and burst together into the same syllable, all syllables of the Brent corpus starting with a stop are marked as starting with a “closure” instead.

In addition, Fig. C1 further illustrates the error distributions for the 25 most frequent phonological syllable boundaries in the EstPhon corpus. The plot shows how boundaries with low recall (e.g., $[i] \rightarrow [m]$; $R = 0.53$ in Table C1) have wide multimodal distributions due to a large number of missed syllable boundaries while high-recall boundaries (e.g., $[k] \rightarrow [k]$) have narrow unimodal distributions. Phone-specific distributions also explain the small separate peak observed in the overall distribution shown in Fig. 5 of the main text: Automatic segmentation at transitions from closure to a stop ($[t:] \rightarrow [t]$ or $[k:] \rightarrow [k]$ at the bottom row the figure) systematically deviate from the manually annotated reference with the model generating these boundaries too early in time. This is not surprising from the acoustic point of view, as the manually annotated boundary between $[k:]$ and $[k]$ is always located at the onset of the burst (release) of $[k]$, whereas the energy and sonority of the signal is zero for the majority of the preceding closure without a single instantaneous minimum. Since the sonority envelope decays to zero after the closure onset, the boundary will become hypothesized to a position where the sonority reaches the noise floor of the process. From a perceptual speech segmentation point of view, this systematic error is hardly relevant since the boundary can be anywhere within the silent closure segment without affecting the spectral cues included in the preceding or following syllable, but for evaluation purposes, it may not count as a correct segmentation.

In addition to the preceding stop boundaries, there are certain boundaries whose detection are slightly delayed with respect to the linguistic reference (e.g., vowel \rightarrow fricative and vowel \rightarrow stop), but are still generally well within the 50-ms evaluation margins.

The basic finding from both corpora is that syllable boundaries with stops (e.g., $/k/$, $/t/$, $/p/$) and fricatives ($/s/$, $/f/$), i.e., sounds with low sonority, on either side of the boundary are the easiest to segment. In contrast, approximants (e.g., $/r/$, $/l/$) and nasals ($/n/$, $/m/$) tend to be involved in more ambiguously syllable boundaries due to larger coarticulatory effects (e.g., nasalization of the neighboring vowels is very common across languages) and generally higher sonority. As can be expected, vowel-to-vowel boundaries have the lowest overall recall due to the high sonority on both sides of the syllable boundary, but are also less frequent in all studied languages.

Table C1

Syllable boundary recalls for EstPhon corpus shown separately for the phones surrounding the syllable boundary. Rows show the preceding phone (last phone of the previous syllable) and columns show the following phone (first phone of the next syllable). Only transitions with 100 or more tokens are shown, covering 66.2% of all syllable boundaries in the corpus. Phones are shown with SAMPA encoding.

	k	k_v	p	p_v	t	t_v	s	s_v	h	h_v	j	v	r	l	l'	A	e	o	m	n	n'	h#	mean														
k	0.98						0.89																	0.93	stops												
p		0.99					0.94																	0.96		stops											
t	0.95	0.91			0.97		0.89				0.46	0.64		0.62					0.54	0.58		0.42		0.70			stops										
t'					0.96																			0.96				stops									
t_v																			0.84	0.87				0.86					stops								
s	0.99	0.97			1.00		0.95				0.95	0.96							0.94	0.92		0.71		0.93	fricatives												
s'					0.98		0.95																	0.97		fricatives											
s:	1.00				0.99		0.95																	0.98			fricatives										
h																			0.71					0.71				fricatives									
h_v																			0.50					0.50					fricatives								
h:	0.98				0.96																	0.56		0.83						fricatives							
r											0.73	0.64												0.68	approx.												
l	0.99				0.98	0.97	0.71							0.75					0.69	0.68				0.82		approx.											
l'										0.73				0.85										0.79			approx.										
l:														0.67										0.67				approx.									
A	0.95	0.92	0.95	0.98	0.94	0.95	0.74	0.95	0.84	0.80	0.38	0.67	0.80	0.65		0.25	0.26	0.34	0.63	0.59		0.49		0.70	vowels												
A_+		0.88																						0.88		vowels											
A:														0.84										0.84			vowels										
A_k	0.94	0.88					0.78																	0.87				vowels									
AA:							0.63				0.69											0.45		0.59					vowels								
e	0.95	0.93	0.92	0.95	0.96	0.92	0.80	0.93		0.74	0.32	0.73	0.80	0.68	0.74		0.19	0.26	0.62	0.59				0.72						vowels							
ee														0.74										0.74							vowels						
e_o														0.75					0.70					0.72								vowels					
i	0.97	0.91	0.84		0.96	0.88	0.81	0.91			0.29	0.71	0.77	0.74			0.12	0.25	0.53	0.55	0.40	0.45		0.65									vowels				
i:							0.73				0.47								0.66					0.62										vowels			
o					0.94		0.76			0.74			0.82	0.62	0.73				0.63	0.58				0.73											vowels		
ɔ																				0.49				0.49												vowels	
u	0.98	0.91	0.98	0.95	0.96	0.89	0.79	0.93			0.31	0.72	0.82	0.53			0.20		0.53	0.44				0.73													vowels
y										0.62				0.70										0.66													
@							0.85																	0.85	vowels												
{	0.95	0.94								0.76			0.69	0.55										0.78		vowels											
m							0.83												0.65					0.74			nasals										
n					0.95	0.91	0.89													0.57				0.83				nasals									
n:					0.94	0.92														0.63				0.83					nasals								
N	0.93	0.90																						0.91						nasals							
N:	0.95	0.86																						0.91							nasals						
h#	0.79		0.78		0.74		0.53				0.15	0.33				0.19	0.17		0.25	0.29				0.42			closures										
t:	0.94				0.75		0.78							0.24					0.36	0.31		0.26		0.52				closures									
k:	0.72						0.83																	0.78					closures								
p:			0.65																					0.65						closures							
mean	0.94	0.90	0.89	0.96	0.94	0.92	0.81	0.93	0.84	0.73	0.50	0.67	0.78	0.66	0.74	0.22	0.19	0.28	0.61	0.58	0.40	0.48										closures					
	stops				fricatives				approximants				vowels			nasals																					

Table C2

Syllable boundary recalls for Brent corpus shown separately for the phones surrounding the syllable boundary. Rows show the preceding phone (last phone of the previous syllable) and columns show the following phone (first phone of the next syllable). Only transitions with 50 or more tokens are shown, covering 56.3% of all syllable boundaries in the corpus. Phones are shown with ARPABET encoding.

	s	f	hh	dh	r	l	w	aa	ae	ah	ao	ax	axr	ay	eh	ey	ih	ix	iy	ow	uw	y	m	n	cl	vcl	mean
k				0.97				0.74	0.80	0.84					0.78	0.90							m	n	cl	vcl	0.84
p					0.75							0.90	0.77														0.80
t	0.78	0.74	0.85	0.92	0.85	0.88	0.82					0.71			0.58						0.72	0.85	0.84	0.79	0.81	0.75	0.79
g								0.82							0.80						0.70						0.77
b											0.82																0.82
d				0.88												0.85				0.75	0.88				0.67	0.74	0.80
s			0.92	0.84											0.76	0.87	0.79							0.80	0.79		0.82
hh						0.83		0.66								0.73											0.74
dh										0.78	0.84																0.81
v				0.80								0.88															0.84
z				0.82						0.85												0.97			0.66		0.83
r			0.68			0.53					0.27											0.31	0.76	0.30	0.45		0.47
l					0.57						0.52	0.51												0.34	0.49		0.49
dx											0.55						0.77	0.66									0.66
ah																						0.36					0.36
ax						0.51																					0.51
ey						0.70																					0.70
ow																	0.25					0.31		0.22			0.26
uh																								0.00			0.00
uw						0.46				0.20												0.45					0.37
m								0.35														0.56				0.72	0.54
n	0.20	0.80	0.66						0.50	0.44									0.64	0.68		0.45		0.34	0.78		0.55
ng																								0.47			0.47
mean	0.49	0.74	0.81	0.84	0.71	0.81	0.64	0.58	0.70	0.65	0.82	0.61	0.83	0.51	0.80	0.71	0.82	0.64	0.66	0.71	0.80	0.56	0.65	0.79	0.44	0.67	
	fricatives				approx.			vowels															nasals		closure		

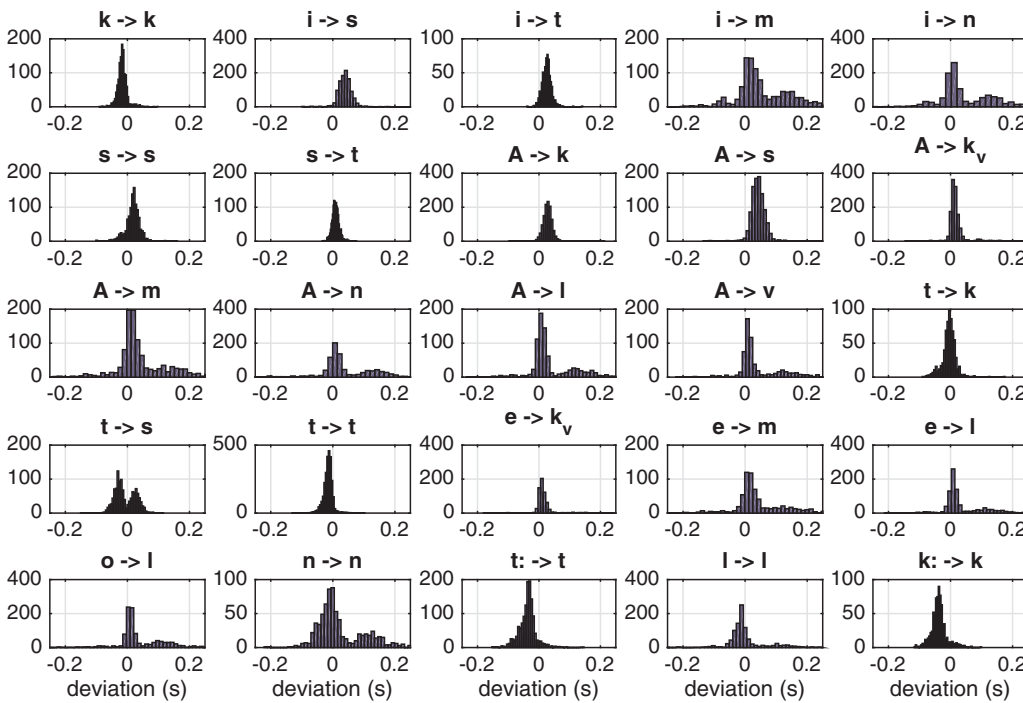


Fig. C1. Distributions for the segmentation accuracies for the 25 most frequent phone transitions across a syllable boundary in the EstPhon corpus. X-axis corresponds to the difference between the hypothesized boundary and the true reference boundary ($t_{\text{hypo}} - t_{\text{ref}}$).

Appendix D. Part of speech analysis

We also investigated how part of speech (POS) categories of the syllables and words were related to their segmentation accuracy. POS tags were available for the Brent and EstPhon corpora and Table D1 shows the full results for the two.

The main finding from Brent IDS speech was that communicatives (e.g., “hey”, “oh”) were most consistently segmented. As for the words, nearly 75% of nouns, 67% of adjectives and 65% of verbs are properly segmented in the Brent data, revealing that the rhythmic structure enables the discovery of a large proportion of content words in their entirety. In contrast, typically shorter words such as auxiliary verbs or prepositions tend to be merged to bigger rhythmic units and have much lower recall as standalone units. Note that the segmentation accuracies for word categories seem to follow the typical order observed in early word learning where infants tend to acquire content words earlier than function words (MacArthur Communicative Development Inventories; Fenson et al., 1993; Frank et al., 2017).

The situation for Estonian words is somewhat different due to fact that large proportion of the vocabulary consists of multisyllabic words. Determiner and communicative words are segmented much more accurately than other words due to their tendency to be more commonly monosyllabic than many other words. Only 56% of noun word tokens and 45% of verbs are correctly segmented while the syllables belonging to same word classes are segmented slightly better (60% and 53%, respectively). In general, the lower overall recall of content words supports the a priori assumption (Section 1.1) that discovery of content words in agglutinative languages such as Estonian and Finnish would greatly benefit from further (e.g., statistical) learning mechanisms. However, it is still somewhat surprising that individual syllables in content words are not segmented any better than in other POS categories. The reason for this is currently unknown. It might be that multisyllabic words of conversational Estonian tend to exhibit large amount of vowel reduction and segmental omissions, a salient characteristic of the closely related Finnish language. This might make the syllabic structure more ambiguous and difficult to annotate than what would be expected on the basis of the predominant CV-CV syllable structure observed in the written language. Also, the adult-directed speaking style in EstPhon might lead to less emphasis on content words than the IDS in Brent data, thereby unifying the clarity of rhythmic structure across POS categories. However, more controlled data with comprehensive annotations would be needed to address IDS/ADS and language comparisons in more detail.

Table D1

Word and syllable unit recalls shown separately for different parts of speech (POS). Only the POS classes with at least 100 word tokens in the data are shown.

EstPhon			Brent				
Words		Syllables	Words		Syllables		
Communicative	0.718	Determiner	0.628	Communicative	0.831	Relative pronoun	0.754
Determiner	0.688	Hesitation	0.621	Relative pronoun	0.806	Communicative	0.661
Hesitation	0.580	Genitive attribute	0.613	Noun	0.748	Determiner	0.613
Adjective	0.568	Noun	0.597	Partitive	0.693	Verb	0.559
Noun	0.556	Adverbial	0.587	Adjective	0.667	Adverbial	0.543
Genitive attribute	0.540	Adjective	0.581	Verb	0.651	Noun	0.528
Adverbial	0.534	Preposition	0.564	Modifier	0.635	Pronoun	0.524
Pronoun	0.511	Pronoun	0.561	Adverbial	0.627	Coordinator	0.513
Preposition	0.473	Communicative	0.554	Determiner	0.588	Preposition	0.495
Conjunction	0.456	Conjunction	0.550	Pronoun	0.580	Modifier	0.484
Verb	0.452	Verb	0.533	Coordinator	0.550	Adjective	0.473
Silence	0.167	Silence	0.013	Quantifier	0.538	Quantifier	0.473
				Preposition	0.518	Copula	0.454
				Copula	0.515	Auxiliary verb	0.394
				Infinitive	0.504	Partitive	0.373
				Auxiliary verb	0.482	Infinitive	0.333

References

- Adriaans, F., & Kager, R. (2010). Adding generalization to statistical learning: The induction of phonotactics from continuous speech. *Journal of Memory and Language*, 62, 311–331.
- Ahissar, E., Nagarajan, S., Ahissar, M., Protopapas, A., Mahncke, H., & Merzenich, M. (2001). Speech comprehension is correlated with temporal response patterns recorded from auditory cortex. *PNAS*, 98(11367–13372), 2001.
- Almpanidis, G., & Kotropoulos, C. (2008). Phonemic segmentation using the generalized Gamma distribution and small sample Bayesian information criterion. *Speech Communication*, 50, 38–55.
- An, G., Brizan, D., & Rosenberg, A. (2013). Detecting laughter and filler pauses using syllable-based features. In *Proceedings of the interspeech-2013* (pp. 178–181). Lyon, France.
- Bertoncini, J., Bijeljac-Babic, R., Jusczyk, P. W., Kennedy, L. J., & Mehler, J. (1988). An investigation of young infants' representation of speech sounds. *Journal of Experimental Psychology*, 117, 21–33.
- Bertoncini, J., & Mehler, J. (1981). Syllables as units in infant speech perception. *Infant Behavior and Development*, 4, 247–260.
- Bijeljac-Babic, R., Bertoncini, J., & Mehler, J. (1993). How do 4-day-old infant categorize multisyllabic utterances? *Developmental Psychology*, 29, 711–721.
- Bortfeld, H., Shaw, K., & Depowski, N. (2013). The miracle year: From basic structure to social communication. In L. Gogate, & G. Hollich (Eds.). *Theoretical and computational models of word learning: Trends in psychology and artificial intelligence*. Hershey PA: IGI Global.
- Brent, M. R., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition*, 81, 31–44.
- Christiansen, M. H., Allen, J., & Seidenberg, M. S. (1998). Learning to segment speech using multiple cues: a connectionist model. *Language and Cognitive Processes*, 13(2), 221–268.
- Clements, G. N. (1990). The role of the sonority cycle in core syllabification. In J. Kingston, & M. E. Beckman (Eds.). *Papers in laboratory phonology 1: Between the grammar and physics of speech* (pp. 283–333). Cambridge: Cambridge University Press.
- Clements, G. N. (2009). Does sonority have a phonetic basis? Comments on the chapter by Bert Vaux. In E. Raimy, & C. Cairns (Eds.). *Contemporary views on architecture and representations in phonology*. MIT Press.
- Cummins, F. (2012). Oscillators and syllables: A cautionary note. *Frontiers in Psychology*. <http://dx.doi.org/10.3389/fpsyg.2012.00364>.
- Cutler, A., Mehler, J., Norris, D., & Segui, J. (1986). The syllable's differing role in the segmentation of French and English. *Journal of Memory and Language*, 25, 385–400.
- Cutler, A., Mehler, J., Norris, D., & Segui, J. (1992). The monolingual nature of speech segmentation by bilinguals. *Cognitive Psychology*, 24, 381–410.
- Cutler, A., & Norris, D. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology*, 14, 113–121.
- Dau, T., Kollmeier, B., & Kohlraus, A. (1997). Modeling auditory processing of amplitude modulation. I. Detection and masking with narrow-band carriers. *Journal of the Acoustical Society of America*, 102, 2892–2905.
- de Saussure, F. (1916). *Cours de linguistique générale*. Paris: Payot.
- Doyle, G., & Levy, R. (2013). Combining multiple information types in Bayesian word segmentation. In *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics* (pp. 117–126). Atlanta, Georgia: Human Language Technologies.
- Dupoux, E. (1993). The time course of prelexical processing: The syllabic hypothesis

- revisited. In G. Altman, & R. Shillcock (Eds.), *Cognitive models of speech processing the second sperlonga meeting* (pp. 81–114). Hove, UK: Lawrence Erlbaum Associates.
- Dutoit, T., Pagel, N., Pierret, F., Bataille, O., & van der Vreken, O. (1996). The MBROLA project: Towards a set of high-quality speech synthesizers free of use for non-commercial purposes. In *Proceedings of the fourth international conference on spoken language processing* (pp. 1393–1396). Philadelphia, PA.
- Eimas, P. (1999). Segmental and syllabic representations in the perception of speech by young infants. *Journal of the Acoustical Society of America*, 105, 1901–1911.
- Esposito, A., & Aversano, G. (2005). Text independent methods for speech segmentation. In G. Chollet, (Ed.), *Lecture notes in computer science: Nonlinear speech modeling* (pp. 261–290). Berlin: Springer Verlag.
- Fenson, L., Dale, P. S., Reznick, J. S., Thal, D., Bates, E., Hartung, J. P., ... Reilly, J. S. (1993). *The MacArthur communicative development inventories: User's guide and technical manual*. Baltimore: Paul H. Brookes Publishing Co.
- Fisher, M. W. (1996). tsylb2. National Institute of Standards and Technology < <http://www.nist.gov/speech/tools/> > .
- Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44, 677–694.
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, 20, 578–585.
- Frank, M. C., Goldwater, S., Griffiths, T. L., & Tenenbaum, J. B. (2010). Modeling human performance in statistical word segmentation. *Cognition*, 117, 107–125.
- Fudge, E. (1969). Syllables. *Journal of Linguistics*, 5, 253–287.
- Galves, A., Garcia, J., Duarte, D., & Galves, C. (2002). Sonority as a basis for rhythmic class discrimination. In *Proceedings of the speech prosody, April 11–13*. France: Aix-en-Provence.
- Gambell, T., & Yang, C. (2006). *Word segmentation: Quick but not dirty*. Manuscript. Yale University.
- Ghitza, O. (2011). Linking speech perception and neurophysiology: Speech decoding guided by cascaded oscillators locked to the input. *Frontiers in Psychology*, 2, 1–13.
- Giraud, A.-L., & Poeppel, D. (2012). Cortical oscillations and speech processing: Emerging computational principles and operations. *Nature Neuroscience*, 15, 511–517.
- Godfrey, J. J., Holliman, E. C., & McDaniel, J. (1992). SWITCHBOARD: telephone speech corpus for research and development. In *Proceedings of the IEEE international conference on acoustics, speech, and signal processing* (pp. 517–520). San Francisco, CA.
- Goldsmith, J. (2011). The syllable. In J. Goldsmith, J. Riggle, & A. Yu (Eds.). *The handbook of phonological theory* (2nd ed.). UK: Wiley-Blackwell Publishing, John Wiley & Sons Ltd.
- Gomez, R. L., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition*, 70, 109–135.
- Grabe, E., & Low, E. L. (2002). Acoustic correlates of rhythm class. In C. Gussenhoven, & N. Warner (Vol. Eds.), *Laboratory phonology: Vol. 7*, (pp. 515–546). Berlin: Mouton de Gruyter.
- Greenberg, S. (1999). Speaking in shorthand – A syllable-centric perspective for understanding pronunciation variation. *Speech Communication*, 29, 159–176.
- Greenberg, S., Carvey, H., Hitchcock, L., & Chang, S. (2003). Temporal properties of spontaneous speech—A syllable-centric perspective. *Journal of Phonetics*, 31, 465–485.
- Gross, J., Hoogenboom, N., Thut, G., Schyns, P., Panzeri, S., Belin, P., & Garrod, S. (2013). Speech rhythms and multiplexed oscillatory sensory coding in the brain. *PLOS Biology*, 11(12), e1001752.
- Guffey, K. (2002). *Spanish syllable structure*. Lanham, MD: University Press of America.
- Hallé, P., & Christia, A. (2012). Global and detailed speech representations in early language acquisition. In S. Fuchs, M. Weirich, D. Pape, & P. Perrier (Eds.). *Speech planning and dynamics*. Frankfurt am Main: Peter Lang.
- Harris, J. (2006). The phonology of being understood: Further arguments against sonority. *Lingua*, 116, 1483–1494.
- Hay, J. F., & Saffran, J. R. (2012). Rhythmic grouping biases constrain infant statistical learning. *Infancy*, 17, 610–641.
- Hooper, J. (1976). *An introduction to natural generative phonology*. New York: Academic Press.
- Hyafil, A., Fontolan, L., Kabdebon, C., Gutkin, B., & Giraud, A.-L. (2015). Speech encoding by coupled cortical theta and gamma oscillations. *eLife*, 4, e06213. <http://dx.doi.org/10.7554/eLife.06213>.
- Jany, C., Gordon, M., Nash, C. M., & Takara, N. (2007). How universal is the sonority hierarchy? A cross-linguistic acoustic study. *Proceedings of the International Congress of Phonetic Sciences*, 1401–1404.
- Jespersen, O. (1920). *Lehrbuch der phonetik* (2nd ed.). Leipzig: Teubner.
- Jones, D. (1960). *An outline of English phonetics*. Cambridge: Cambridge University Press.
- Jusczyk, P. W. (1993). From general to language specific capacities: The WRAPSA model of how speech perception develops. *Journal of Phonetics*, 21, 3–28.
- Jusczyk, P. W., Bertoncini, J., Bijeljac-Babic, R., Kennedy, L. J., & Mehler, J. (1990). The role of attention in speech perception by infants. *Cognitive Development*, 5, 265–286.
- Jusczyk, P. W., & Derrah, C. (1987). Representation of speech sounds by young infants. *Developmental Psychology*, 23, 648–654.
- Jusczyk, P. W., Houston, D. M., & Newsome, M. (1999). The beginnings of word segmentation in English-learning infants. *Cognitive Psychology*, 39, 159–207.
- Jusczyk, P. W., Kennedy, L. J., & Jusczyk, A. M. (1995). Young infants retention of information about syllables. *Infant Behavior and Development*, 18, 27–41.
- Jusczyk, P. W., & Luce, P. A. (2002). Speech perception and spoken word recognition: past and present. *Ear & Hearing*, 23(1), 2–40.
- Jusczyk, P. W., & Thompson, E. (1978). Perception of a phonetic contrast in multisyllabic utterances by 2-month-old infants. *Perception & Psychophysics*, 23, 105–109.
- Kahn, D. (1976). *Syllable based generalizations in English phonology*, Ph.D. dissertation, Department of Linguistics and Philosophy, MIT, Cambridge.
- Krakow, R. A. (1999). Physiological organization of syllables: A review. *Journal of Phonetics*, 27, 23–54.
- Kvale, K. (1993). *Segmentation and labelling of speech*. Doctoral thesis, The Norwegian Institute of Technology, Trondheim, Norway.
- Ladefoged, P. (2000). *A course in phonetics* (4th ed.). Boston, MA: Thomson Wadsworth.
- Lee, C. S., & Todd, N. P. (2004). Towards an auditory account of speech rhythm: Application of a model of the auditory 'primal sketch' to two multi-language corpora. *Cognition*, 93, 225–254.
- Leoni, F. A. (2015). The boundaries of the syllable. In D. Russo (Ed.). *The notion of syllable across history: Theories and analysis*. UK: Cambridge Scholars Publishing.
- Lenneberg, M. (2009). Segmental features in spontaneous and read-aloud Finnish. In V. de Silva, & R. Ullakonjoja (Eds.). *Phonetics of Russian and Finnish* (pp. 145–166). Frankfurt am Main: Peter Lang.
- Leong, V. (2012). *Prosodic rhythm in the speech amplitude envelope: Amplitude modulation phase hierarchies (AMPHs) and AMPH models*, Doctoral dissertation. University of Cambridge.
- Leong, V., & Goswami, U. (2014). Assessment of rhythmic entrainment at multiple timescales in dyslexia: Evidence for disruption to syllable timing. *Hearing Research*, 308, 141–161.
- Leong, V., & Goswami, U. (2015). Acoustic-emergent phonology in the amplitude envelope of child-directed speech. *PLoS ONE*. <http://dx.doi.org/10.1371/journal.pone.0144411>.
- Leong, V., Stone, M. A., Turner, R. E., & Goswami, U. (2014). A role for amplitude modulation phase relationships in speech rhythm perception. *The Journal of the Acoustical Society of America*, 136, 366–381.
- Liberman, I. Y., Shankweiler, D., Fischer, W. F., & Carter, B. (1974). Explicit syllable and phoneme segmentation in the young child. *Journal of Experimental Child Psychology*, 18, 201–212.
- Lippus, P., Tuisk, T., Salvete, N., & Teras, P. (2013). *Phonetic corpus of Estonian spontaneous speech*. Institute of Estonian and General Linguistics, University of Tartu <https://doi.org/10.15155/TY.0000>.
- Luo, H., & Poeppel, D. (2007). Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron*, 54, 1001–1010.
- Malmberg, B. (1963). *Phonetics*. New York: Dover Publications Inc.
- Marchand, Y., Adsett, C. R., & Damper, R. I. (2007). Evaluating automatic syllabification algorithms for English. In *Proceedings of the 6th international speech communication association (ISCA) workshop on speech synthesis* (pp. 316–321), Bonn, Germany.
- Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, 283, 77–80.
- Mehler, J., & Christophe, A. (1995). Maturation and learning of language in the first year of life. In M. Gazzaniga (Ed.). *The cognitive neurosciences* (pp. 943–954). Cambridge, MA: MIT Press.
- Mehler, J., Dommergues, J. Y., Frauenfelder, U., & Segui, J. (1981). The syllable's role in speech segmentation. *Journal of Verbal Learning and Verbal Behavior*, 20, 298–305.
- Mehler, J., Dupoux, E., & Segui, J. (1990). Constraining models of lexical access: The onset of word recognition. In G. Altman (Ed.). *Cognitive models of speech processing* (pp. 236–262). Cambridge Mass: MIT Press.
- Mehler, J., Jusczyk, P., Lambert, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. *Cognition*, 29, 143–178.
- Mermelstein, P. (1975). Automatic segmentation of speech into syllabic units. *Journal of the Acoustical Society of America*, 58, 880–883.
- Meylan, S., Kurumada, S., Börschinger, B., Johnson, M., & Frank, M. C. (2012). Modeling online word segmentation performance in structured artificial languages. In N. Miyake, D. Peebles, & R. Cooper (Eds.), *Proceedings of the 34th annual meeting of the cognitive science society* (pp. 2002–2007). Sapporo, Japan.
- Morais, J., Bertelson, P., & Alegria, J. (1986). Literacy training and speech segmentation. *Cognition*, 24(1–2), 45–64.
- Morais, J., Content, A., Cary, L., Mehler, J., & Segui, J. (1989). Syllabic segmentation and literacy. *Language and Cognitive Processes*, 4(1), 56–67.
- Morgan, N., & Fosler-Lussier, E. (1998). Combining multiple estimators of speaking rate. In *Proceedings of the IEEE international conference on acoustics, speech and signal processing (ICASSP-98)* (pp. 729–732). Seattle, WA.
- Nasukawa, K. (2007). Relational properties in phonology: precedence and dependency. In *CUNY conference on precedence relations*.
- Nazzi, T., Bertoncini, J., & Mehler, J. (1998). Language discrimination by newborns: Toward an understanding of the role of rhythm. *Journal of Experimental Psychology*, 24(3), 756–766.
- Nazzi, T., Iakimova, G., Bertoncini, J., Frédonie, S., & Alcantara, C. (2006). Early segmentation of fluent speech by infants acquiring French: Emerging evidence for crosslinguistic differences. *Journal of Memory and Language*, 54, 283–299.
- Nespor, M., Shukla, M., & Mehler, J. (2011). Stress-timed vs. syllable-timed languages. In M. van Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.). *The Blackwell companion to phonology*. West Sussex, UK: John Wiley & Sons Ltd.
- Newport, E., & Aslin, R. (2004). Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, 48, 127–162.
- Ngon, C., Martin, A., Dupoux, E., Cabrol, D., Dutat, M., & Peperkamp, S. (2013). (Non) words, (non)words, (non)words: Evidence for a protolexicon during the first year of life. *Developmental Science*, 16, 24–34.
- Nishibayashi, L.-L., Goyet, L., & Nazzi, T. (2015). Early speech segmentation in French-learning infants: Monosyllabic words versus embedded syllables. *Language and Speech*, 58, 334–350.
- Nusbaum, H. C., & DeGroot, J. (1991). The role of syllables in speech perception. In M. S. Ziolkowski, M. Noske, & K. Deaton (Eds.). *Papers from the parasession on the syllable in phonetics and phonology*. Chicago: Chicago Linguistic Society.
- Obin, N., Lamare, F., & Roebel, A. (2013). Syll-O-Matic: An adaptive time-frequency representation for the automatic segmentation of speech into syllables. In *Proceedings of*

- the IEEE international conference on acoustics, speech and signal processing (ICASSP-2013) (pp. 6699–6703). Vancouver, BC.
- Ohala, J. (1990). There is no interface between phonology and phonetics: A personal view. *Journal of Phonetics*, 18, 153–171.
- Ohala, J. (1992). *Alternatives to the sonority hierarchy for explaining segmental sequential constraints. Papers from the parasession on the syllable*. Chicago: Chicago Linguistics Society 319–338.
- Oppenheim, A. V., Willsky, A. S., & Young, I. T. (1983). *Signals and systems*. New Jersey: Prentice Hall.
- Parker, S. G. (2002). *Quantifying the sonority hierarchy*. PhD dissertation. Graduate School of the University of Massachusetts Amherst, MA.
- Patterson, R. D., Robinson, K., Holdsworth, J., McKeown, D., Zhang, C., & Allerhand M. (1992). Complex sounds and auditory images. In Y. Cazals, L. Demany, K. Horner (Eds.), *Proceedings of the 9th international symposium on hearing, auditory physiology and perception* (pp. 123–177). Pergamon, Oxford.
- Peelle, J. E., Gross, J., & Davis, M. H. (2013). Phase-locked responses to speech in human auditory cortex are enhanced during comprehension. *Cerebral Cortex*, 23, 1378–1387.
- Perruchet, P., & Tillman, B. (2010). Exploiting multiple sources of information in learning and artificial language: Human data and modeling. *Cognitive Science*, 34, 255–285.
- Perruchet, P., & Vinter, A. (1998). PARSE: A model for word segmentation. *Journal of Memory and Language*, 39, 246–263.
- Phillips, L., & Pearl, L. (2012). 'Less is More' in Bayesian word segmentation: When cognitively plausible learners outperform the ideal. In N. Miyake, D. Peebles, & R. Cooper (Eds.), *Proceedings of the 34th annual meeting of the cognitive science society* (pp. 863–868), Sapporo, Japan.
- Port, R. F., & Leary, A. P. (2005). Against formal phonology. *Language*, 81, 927–964.
- Price, P. J. (1980). Sonority and syllabicity: Acoustic correlates of perception. *Phonetica*, 37, 327–343.
- Prince, A., & Smolensky, P. (1993). *Optimality theory: Constraint interaction in generative grammar*. Rutgers University Center for Cognitive Science Technical Report no. 2, New Brunswick, NJ.
- Räsänen, O. (2011). A computational model of word segmentation from continuous speech using transitional probabilities of atomic acoustic events. *Cognition*, 120, 149–176.
- Räsänen, O. (2014). Basic cuts revisited: Temporal segmentation of speech into phone-like units with statistical learning at a pre-linguistic level. In *Proceedings of the 36th annual conference of the cognitive science society* (pp. 2817–2822), Quebec, Canada.
- Räsänen, O., Laine, U. K., & Altsosaar, T. (2009). An improved speech segmentation quality measure: The R-value. In *Proceedings of the interspeech-09* (pp. 1851–1854), Brighton, England.
- Räsänen, O., & Laine, U. K. (2013). Time-frequency integration characteristics of hearing are optimized for perception of speech-like acoustic patterns. *Journal of the Acoustical Society of America*, 134(407–419), 2013.
- Räsänen, O., & Rasilo, H. (2015). A joint model of word segmentation and meaning acquisition through cross-situational learning. *Psychological Review*, 122(4), 792–829.
- Redford, M. A., & Randall, P. (2005). The role of juncture cues and phonological knowledge in English syllabification judgments. *Journal of Phonetics*, 33, 27–46.
- Rosenberg, A. (2010). AuToBI - A Tool for Automatic ToBI Annotation. In *Proceedings of the interspeech-2010* (pp. 146–149). Makuhari, Japan.
- Rytting, A., Brew, C., & Fosler-Lussier, E. (2010). Segmenting words from natural speech: Subsegmental variation in segmental cues. *Journal of Child Language*, 37(3), 513–543.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70, 27–52.
- Scharenborg, O., Ernestus, M., & Wan, V. (2007). Segmentation of speech: Child's play? In *Proceedings of the interspeech-07* (pp. 1953–1956). Antwerp, Belgium.
- Segui, J., Frauenfelder, U., & Mehler, J. (1981). Phoneme monitoring, syllable monitoring and lexical access. *British Journal of Psychology*, 72, 471–477.
- Seidl, A., & Johnson, E. (2006). Infant word segmentation revisited: Edge alignment facilitates target extraction. *Developmental Science*, 9, 565–573.
- Shukla, M., Nespor, M., & Mehler, J. (2007). An interaction between prosody and statistics in the segmentation of fluent speech. *Cognitive Psychology*, 54, 1–32.
- Shukla, M., White, K. S., & Aslin, R. N. (2011). Prosody guides the rapid mapping of auditory word forms onto visual objects in 6-month-old infants. *PNAS*, 108(15), 6038–6043.
- Stetson, R. H. (1951). *Motor phonetics: A study of speech movements in articulation* (2nd ed.). Amsterdam: North Holland.
- Suomi, K., Toivanen, J., & Ylitalo, R. (2008). *Finnish sound structure*. Oulu, Finland: Oulu University Press.
- Swingle, D. (2005). Statistical clustering and the contents of the infant vocabulary. *Cognitive Psychology*, 50, 86–132.
- Tesar, B., & Smolensky, P. (1998). Learnability in optimality theory. *Linguistic Inquiry*, 29, 229–268.
- Thiessen, E. D., & Saffran, J. R. (2003). When cues collide: Use of stress and statistical cues to word boundaries by 7- and 9-month-old infants. *Developmental Psychology*, 39(4), 706–716.
- van der Hulst, H. (2005). The molecular structure of phonological segments. In P. Carr, J. Durand, & C. Ewen (Eds.), *Headhood, elements, specification and contrastivity* (pp. 193–234). Amsterdam: John Benjamins.
- Viemeister, N. F. (1979). Temporal modulation transfer functions based upon modulation thresholds. *Journal of the Acoustical Society of America*, 66, 1364–1380.
- Villing, R., Timoney, J., Ward, T., & Costello, J. (2004). Automatic blind syllable segmentation for continuous speech. In *Proceedings of the Irish signals and systems conference (ISSC 2004)*. Belfast, Northern Ireland.
- Villing, R., Ward, T., & Timoney, J. (2006). Performance limits for envelope-based automatic syllable segmentation. In *Proceedings of the ISSC-2006* (pp. 521–526). Dublin, Ireland.
- Wagner, P. (2008). The rhythm of language and speech: Constraining factors, models, metrics and applications. Habilitation thesis, University of Bonn, Germany.
- Wang, D., & Narayanan, S. (2007). Robust speech rate estimation for spontaneous speech. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(8), 2190–2201.
- Whitney, W. D. (1874). *Oriental and linguistic studies*. Second Series. New York: Scribner, Armstrong & Co.
- Ziegler, J. C., & Goswami, U. (2005). Reading acquisition, developmental dyslexia, and skilled reading across languages: A psycholinguistic grain size theory. *Psychological Bulletin*, 131, 3–29.