Cognition 166 (2017) 298–313

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Changing word usage predicts changing word durations in New Zealand English

Márton Sóskuthy [a,*], Jennifer Hay [b]

[a] University of York, United Kingdom
[b] University of Canterbury, New Zealand

ARTICLE INFO

Article history:
Received 22 January 2016
Revised 17 May 2017
Accepted 19 May 2017
Available online 6 June 2017

Keywords:
New Zealand English
Word duration
Word frequency
Informativity
Production-perception loop
Language change

ABSTRACT

This paper investigates the emergence of lexicalized effects of word usage on word duration by looking at parallel changes in usage and duration over 130 years in New Zealand English. Previous research has found that frequent words are shorter, informative words are longer, and words in utterance-final position are also longer. It has also been argued that some of these patterns are not simply online adjustments, but are incorporated into lexical representations. While these studies tend to focus on the synchronic aspects of such patterns, our corpus shows that word-usage patterns and word durations are not static over time. Many words change in duration and also change with respect to frequency, informativity and likelihood of occurring utterance-finally. Analysis of changing word durations over this time period shows substantial patterns of co-adaptation between word usage and word durations. Words that are increasing in frequency are becoming shorter. Words that are increasing/decreasing in informativity show a change in the same direction in duration (e.g. increasing informativity is associated with increasing duration). And words that are increasingly appearing utterance-finally are lengthening. These effects exist independently of the local effects of the predictors. For example, words that are increasing utterance-finally lengthen in all positions, including utterance-medially. We show that these results are compatible with a number of different views about lexical representations, but they cannot be explained without reference to a production-perception loop that allows speakers to update their representations dynamically on the basis of their experience.

1. Introduction

It is well-established that a number of usage factors affect word duration – including frequency, the word's predictability in context, and the position of the word in relation to utterance boundaries. In theory, there are two ways in which such effects can be realized (see, e.g. Bybee, 2002; Jaeger & Buz, in press). First, they can manifest as context-dependent, local adjustments that apply online during speech production. The existence of such local effects is uncontroversial, and they are the main focus of a large portion of the literature on variation in word duration. But usage-based effects can also manifest as offline lexicalized changes that affect words regardless of their context. Recent research based on synchronic corpus data shows that such lexical effects may exist alongside local effects (Seyfarth, 2014), and suggests that the two are linked: changes to lexical representations arise through repeated exposure to local effects (this idea is already anticipated in Paul, 1880, p. 46).

This paper presents an empirical investigation of the emergence of lexical effects on word duration. Such lexical effects likely exist at all points in the history of a language, so it is not possible to look at their 'ultimate' origin. Instead, we focus on a specific question that can be investigated using relatively recent historical data: what happens to word duration when a word's usage patterns are not stable over time? In such situations, it should be possible to directly observe the emergence of lexical effects in the form of co-adaptation between usage and form. Therefore, we ask the following questions: Can patterns of changing word usage predict patterns of changing word production? Is there evidence that lexical representations are directly impacted by changing word usage patterns?

The research presented in this paper extends previous work substantially by tracking word duration trajectories and changes to word usage over time in a diachronic corpus. Our data set comes

* Corresponding author at: Department of Language and Linguistic Science, University of York, Heslington, York YO10 5DD, United Kingdom.
E-mail address: marton.soskuthy@york.ac.uk (M. Sóskuthy).

http://dx.doi.org/10.1016/j.cognition.2017.05.032
0010-0277/

from the spoken Origins of New Zealand English corpus (ONZE, Gordon, Maclagan, & Hay, 2007), which contains speech samples from over 500 speakers born between 1851 and 1987. We track changes to 698 content words represented by more than 270,000 tokens, focusing on word usage, word duration and the extent to which they change together.

Using this unique data set, we are able to obtain a direct view of the accumulation of usage-based effects in lexical representations over time. These show up in the form of robust parallels between changes in word duration and usage. We suggest that these findings are compatible with a range of different views about lexical representations, but are difficult to explain without reference to the so-called *production-perception loop* (Pierrehumbert, 2001; Wedel, 2007).

The paper is structured as follows. In Section 2, we summarize observations about patterns of variation in word duration and briefly describe the potential pathways that can lead to lexical patterns, with special emphasis on the production-perception loop. Section 3 sets out our synchronic and diachronic hypotheses relying on the discussion in the preceding section. Section 4 first gives an overview of the spoken diachronic corpus that serves as the basis of the project, and then defines our key variables. Section 5 plays a mainly descriptive role, presenting general patterns of change in word duration and usage factors based on the corpus, and setting the scene for the main statistical analysis presented in Section 6. Section 7 concludes the paper with a discussion of the results and an evaluation of the hypotheses, along with some more general conclusions about the nature of language change.

2. Background

2.1. Word duration and usage factors

One of our key variables is word duration, defined as the duration of spoken word forms measured in seconds. Word duration varies substantially as a function of frequency, predictability, repetition, syntactic probability and a range of other variables (Bell et al., 2003; Bell, Brenier, Gregory, Girand, & Jurafsky, 2009; Gahl, 2008; Jurafsky, Bell, Gregory, & Raymond, 2001; Seyfarth, 2014; Tily et al., 2009; Whalen, 1991). This paper uses the term *usage factor* to refer to these variables collectively. We do not assign special theoretical significance to variation in word duration, and simply take it to be one of the many phonetic reflexes of more general processes of hypo- and hyper-articulation (cf. Lindblom, Guion, Hura, Moon, & Willerman, 1995) conditioned by usage factors. Other examples of such reflexes include variation in segmental and syllabic duration, the peripherality of vowels and consonant deletion (Aylett & Turk, 2006; Bybee, 2002; Cohen Priva, 2015; Jurafsky et al., 2001).

As noted in the introduction, patterns of variation in word duration can be divided into two types based on the way they are expressed: as differences between tokens of the same word in different local contexts, or as context-independent differences across multiple lexical items. These patterns will be labelled *local* and *lexical*, respectively.

An example of a local pattern is the effect of predictability from the preceding or following context: words tend to be shorter in predictable contexts (Bell et al., 2009; Jurafsky et al., 2001; Seyfarth, 2014). Since words typically appear both in high and low predictability contexts (e.g. the word *hunt* in *witch hunt* vs. *which hunt*; Lieberman, 1963), they display within-item local variation based on predictability. There are a variety of proposals about the mechanisms through which contextual predictability comes to be related to reduced forms, some of which relate to speaker-based factors such as ease of access or planning, and some of which are more listener-oriented, relating to appropriately conveying the intended message. A good recent outline of various accounts is provided in Jaeger and Buz (in press). The topic of interest in this paper is the potential accumulated consequences of these local forces at the lexical level.

As opposed to local patterns, lexical patterns are stable across contexts for each word, but vary across different words. A simple example of a lexical pattern is the effect of unigram word frequency: high-frequency items tend to be shorter than low-frequency items (e.g. Bell et al., 2009; Gahl, 2008). Since the unigram frequency of a word is not context-dependent, a given lexical item will always show the same effect of frequency, and the effect of frequency can only be seen by comparing multiple lexical items.

Before we take a closer look at the specific usage factors investigated in this paper, it will be useful to provide a brief overview of the types of correlations we may observe between changes in usage factors and word durations. There is a trivial sense in which shifts in the distribution of local conditioning factors may lead to changes in observed word durations. All things being equal, a word that becomes more predictable in a given context will also undergo more shortening in that context, which also lowers its average duration. Such parallel changes between word duration and usage factors are superficial in the sense that they do not affect lexical representations. Although the surface distribution of word durations may change along with the word's predictability in specific contexts, this change simply and directly reflects the online local reductive forces at work. Those tokens of the word that happen to occur in low-predictability contexts will not undergo shortening.

In this paper, we are particularly interested in lexical changes that go beyond local conditioning factors and whose effects are not dependent on the immediately local context – in other words, changes that arguably take place at the level of lexical representations. An example of such an effect is presented by Seyfarth (2014), who demonstrates that words which tend to occur in predictable contexts are shorter *even when their local predictability is low*. He argues that such lexical effects reflect stored patterns of reduction, which come about through repeated exposure to local reductive biases. The crucial step in his analysis is the separation of two different effects: local predictability and a cumulative measure of predictability calculated over all contexts for a given word, called informativity (cf. below). He shows that informativity has an independent contribution to word duration even after local predictability and a range of other control variables have been taken into account. This paper follows Seyfarth (2014) in separating local and lexical measures and looking for an independent contribution of the latter in an attempt to detect changes that affect lexical representations.

We focus on three main groups of usage factors: predictability, position within the utterance and frequency. Predictability can be defined on many different levels and in many different ways. One of the simplest definitions is based on immediately adjacent words: the conditional probability of a word x given a preceding or a following word y, which is usually approximated through the following equation (where $p(x|y)$ stands for the conditional probability of x given y, $c(xy)$ is the number of times x and y occur together in a corpus and $c(y)$ is the frequency of y in the same corpus; see e.g. Bell et al., 2009; Jurafsky et al., 2001; Seyfarth, 2014):

$$p(x|y) = \frac{c(xy)}{c(y)} \quad (1)$$

As explained above, predictability is a local measure, with a corresponding lexical measure called informativity. Informativity is closely related (although not identical) to average predictability.