The Effect of Phonological Parallelism in Coordination: Evidence from Eye-tracking

Amit Dubey and Frank Keller and Patrick Sturt
Human Communication Research Centre, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW, UK
{amit.dubey,patrick.sturt,frank.keller}@ed.ac.uk

Abstract

In this paper we report an eye-tracking experiment designed to investigate syntactic and phonological parallelism effects in comprehension. Eye-movements were recorded while participants read sentences that contained particle verb constructions. Each experimental item included a coordinated verb phrase (VP), whose two conjuncts either exhibited parallel syntactic forms in terms of particle placement (e.g., *The lawyer won over the jury and fought off the developers*), or did not (e.g., *The lawyer won the jury over and fought off the developers*). In addition to the manipulation of syntactic form, the number of syllables intervening between the verb and the particle ranged from zero to five across the item set. Linear regression analysis revealed a reading time advantage for VPs that were parallel in terms of the number of intervening syllables. However, processing was not speeded by parallelism of syntactic form per se. We argue that theories that seek to explain syntactic parallelism need to take account of phonological length.

Introduction

Research in psycholinguistics has shown evidence for a general facilitation for the processing of structure that is structurally similar to recently processed material. For example, work on *syntactic priming* in production (Bock, 1986; Pickering & Branigan, 1998) shows that people prefer to produce sentences using syntactic structures that have recently been produced. Although most of the work in syntactic priming concerns language production, there are some studies that show a similar facilitation in comprehension. For example, Branigan et al. (1995) showed that whole-sentence reading times for garden path sentences were reduced when a similar garden path sentence had been read in an immediately preceding trial, while Branigan et al. (2005) showed that picture-matching latencies were facilitated when the relevant picture matched the syntactic structure of the preceding trial.

A phenomenon that is closely related to syntactic priming in comprehension is the *parallelism effect*. In an eye-tracking experiment investigating noun phrase coordination, Frazier et al. (2000) showed that the second conjunct *a short poem* was read more quickly in (1-a), when it matched the form of the first conjunct, than in (1-b), when it did not.

(1) a. Terry wrote a long novel and *a short poem* during her sabbatical
    b. Terry wrote a novel and *a short poem* during her sabbatical

Frazier et al. interpreted the parallelism effect as a syntactic phenomenon—in other words, the relative facilitation in (1-a) was due to the fact that both noun phrase conjuncts had the form DETERMINER ADJECTIVE NOUN. However, there are other possible interpretations of the effect. One possible contributing factor is the relative lengths of the two noun phrase conjuncts. For example, the two conjuncts in (1-a) have the same number of syllables, while those in (1-b) do not (henceforth, we will assume that the syllable is the appropriate measure for length effects, and we will refer to such effects in terms of phonology). Another potentially relevant difference is that the determiner and the head noun are separated by one intervening syllable for both conjuncts in (1-a), while this is not the case in (1-b).

Parallelism has also been studied in corpus-based investigations. Dubey et al. (2005) show that the frequencies of various noun phrase structures are increased when the same structure has already appeared in the preceding text. Dubey et al. (2005) also show that the effect can be found whether or not the noun phrases in question are coordinated, although the effect is stronger in coordinated environments. Recent corpus work (Gries, 2005; Szmrecsanyi, 2005; Reitter et al., 2006) has shown that the phenomenon is general across a range of construction types in corpora.

The experiment which we report here investigates the parallelism phenomenon in English particle verbs (also known as phrasal verbs). Particle verbs are of particular interest for the investigation of parallelism, because they allow for a syntactic alternation which has only a minimal effect on meaning. Consider (2-a) and (2-b), for example:

(2) a. The lawyer won over the jury.
    b. The lawyer won the jury over.

Besides information structure differences, these two sentences have identical meanings, while they differ in syntactic form.

A second reason for using particle verbs in our study is that the alternation allows us to compare a purely syntactic view of parallelism with a view that takes account of phonological factors. Consider the following:

(3) a. The lawyer won over the jury.
    b. The campaigner paid bribes out.
    c. Graham knocked the shed down.
    d. The lawyer won the jury over.
    e. The politician handed the estimate in.
    f. The lawyer fought the developers off.

In the sequence of sentences in (3), the number of syllables intervening between the particle and the verb increases from

Proceedings of the 2nd European Cognitive Science Conference, 131–136. Delphi, 2007.

zero in (3-a) to five in (3-f). Thus, from a phonological point of view, the degree of separation can be seen as varying quasi-continuously from (3-a) to (3-f). This contrasts with a syntactic view. In terms of the syntactic rules (3-a) differs from all the other sentences, because (3-a) uses a VERB NP PARTICLE template, while all the other examples use a VERB PARTICLE NP template.

In the theoretical linguistics and psycholinguistics literature, it has been noted that particle verb sentences become intuitively more difficult as the degree of separation between the verb and particle increases (Hawkins, 1995; Frazier & Fodor, 1978).[1] More recently, Gries (1999) has argued that this effect of constituent length can be subsumed under a more general theory in which processing difficulty is explained by semantic, pragmatic and discourse-level factors. However, to our knowledge, this phenomenon has not been studied using on-line techniques. Therefore, a third objective of our paper is to examine how processing difficulty is affected by the degree of separation between verb and particle. This in turn allows us to consider the question of whether the degree of the parallelism preference is affected by the degree of processing difficulty involved in the relevant structure.

Experiment

This experiment had a dual purpose: Firstly, we wanted to test whether the parallelism effect demonstrated by Frazier et al. (2000) for NP coordination generalizes to VP coordination. The second aim was to determine whether phonological factor influence parallelism, such as the number of syllables intervening between a verb and its particle in particle verb constructions such as the ones in (2).

We follow previous work on parallelism in terms of methodology: experimental participants had to read written stimuli while their eye-movements were recorded using an eye-tracker. The eye-tracking record allows us to measure the predicted speedup effects in the second conjunct with great accuracy.

Method

Participants Thirty-two native English speakers (students from the University of Edinburgh) took part in this study, receiving £4 subject payment. Participants were tested in individual sessions, each of which took about 25 minutes to complete.

Materials We designed 28 different materials, each consisting of a sentence with VP coordination involving particle verbs in both conjuncts (no particle verb was used in more than one material). Four different conditions were generated for each of the materials by varying whether the particle was adjacent to the verb or moved in either the first or the second VP. An example material in all four conditions is as follows:

(4) a. Before the lawyer won over the jury and fought off the developers, the project was stalled.
b. Before the lawyer won the jury over and fought off the developers, the project was stalled.
c. Before the lawyer won over the jury and fought the developers off, the project was stalled.
d. Before the lawyer won the jury over and fought the developers off, the project was stalled.

The items were allocated to four lists. Each list contained each of the 28 material in one of the four conditions (according to a Latin square). A list of 90 filler sentences was also generated. Eight participants each read one of the lists (assigned at random), as well as all of the fillers. Items and fillers were presented in random sequence generated for each participant, preceded by five practice items.

Procedure Participants were seated approximately 65 cm from a 21” color monitor with 1024 $\times$ 768 pixel resolution; twenty-four pixels equaled about one degree of visual angle. Participants wore an SR Research Eyelink II head-mounted eye-tracker running at 500 Hz sampling rate. Viewing was binocular, but only the participant’s dominant eye was tracked (the right eye for about 68% of the participants, as determined by a simple parallax test prior to the experiment). Participants were instructed to avoid strong head movements throughout the experiment. A USB gamepad was used to record button responses. Stimulus presentation and data recording were controlled by two PCs running EyeTrack, experimental software developed at the University of Massachusetts.

At the start of the experiment, the experimenter performed the standard Eyelink calibration routine, which involves participants looking at a grid of nine fixation targets in random succession. Then a validation phase followed to test the accuracy of the calibration against the same targets. Calibration and validation was repeated at least twice throughout the experiment, or if the experimenter noticed that measurement accuracy was poor (e.g., after strong head movements or a change in the participant’s posture).

Each trial was structured as follows: first a gaze trigger was displayed on the left of the screen. Once the participants had fixated the gaze trigger, the trial began, with the trigger being replaced by the first letter of the sentence. Once the participant had finished reading the sentences, he/she pressed a button to signal the end of the trial. Participants were instructed to read the sentences attentively, so that they were able to answer subsequent questions. In 25% of the cases, the trial was followed by a question on the screen, replacing the sentences. Whenever such a question appeared, subjects had to answer it by pressing either the “yes” button or the “no” button on the gamepad.

Results

Data Analysis The raw eye-tracking data was processed with EyeDoctor, part of the UMass suite of eye-tracking tools. EyeDoctor was used to correct cases of vertical drift of the tracker, and to delete abandoned trials, fillers, and practice items. Fixations shorter than 80 ms (approximately 2–3% of all fixations) were pooled with preceding or following fixations. Times for blinks were added to the immediately preceding fixations (assuming that processing does not pause during a blink) and fixations outside the screen area (less than 1% of all fixations) were deleted.

[1] Although we know of no result applying Gibson’s Locality Theory (Gibson, 1998) to processing particle verb constructions, it is nonetheless interesting to note this theory may provide some explanation as to why non-adjacent particle verb constructions are more difficult to process.

The data was then analyzed with UMass EyeDry, which computes a range of eye-tracking measures based on a region definition provided by the experimenter. We defined four regions as follows:

- Region r_1: all material up to end of the subject NP (*Before the lawyer* in (4-a));
- Region r_2: first VP (*won over the jury* in (4-a));
- Region r_3: conjunction *and*;
- Region r_4: second VP (*fought the developers off* in (4-a));
- Region r_5: all remaining material (*the project was stalled* in (4-a)).

Region r_4 is the critical region on which we expect to see the parallelism effect. Region r_5 is the spillover region, and might also show an effect. On for these regions, we computed the following reading measures: first-pass time, number of regressions in, number of regressions out, regression path time, second pass time, and total time.

The *first-pass reading times* are the sum of the durations of all initial fixations, from the time that the eye gaze first enters the region from the left until the gaze moves outside the region, either to left or right. *Regression path times* are the sum of all fixations that are made (including to the left of the region) before the eye gaze first moves to the right of the word. *Second pass reading times time* are the sum of all fixations that are made on the region after the region has been fixated and exited for the first time. Finally, *Total time* is the sum of all fixations on the region, and equates to the sum of first-pass and second-pass reading times.

We calculated two measures based on the frequency of regressive saccades: *regressions out* refers to the percentage of trials in which the reader makes a regressive saccade to the left of the region, before any fixations are made to the right of the region. Finally *regressions in* refers to the percentage of trials in which the region was re-fixated following a regressive saccade from the right.

We also coded the experimental stimuli as follows. A categorical variable was introduced that indicates whether the verb and its particle are contiguous (cont) or non-contiguous (non) in both VPs. We will refer to these as vp1 and vp2, respectively. To enable more fine-grained analysis of particle movement effects, we coded also the number of syllables intervening between verb and particle in the first VP and the second VP. We will refer to these variables as syl1 and syl2. Both variables ranges from 0 to 5.

Phonological Parallelism As noted above, we measure the parallelism effect using two variables, syl1, which indicates the number of syllables between the verb and the particle in the first VP, and syl2, which is the same measure for VP_2.

To analyze the data, we regress reading time against the syl1 and syl2 variables. Noting that a verb and particle are more likely to be adjacent when the phrasal verb takes a non-pronominal NP object, we would expect a local minimum when $\mathsf{syl2} = 0$. Under the hypothesis that syllabic parallelism influences reading times, we would further expect a local minimum as the difference between syl1 and syl2 increase. This will be encoded in a new variable $\mathsf{absdiff} = abs(\mathsf{syl1} - \mathsf{syl2})$. As we expect these variables to affect reading times linearly, we use linear regression. The regression model used was

$$r_4 \sim \mathsf{syl1} * \mathsf{syl2} + \mathsf{absdiff}$$

Note that the $*$ represents all possible interactions and main effects between the 'multiplied' variables. Regressing syl1, syl2 and absdiff against reading times produced no statistically significant results ($P > .1$ for all coefficients) for first pass, regression path, regressions in, regressions out, and second path reading times.

This regression model, though, is unable to capture an interesting property of particle verbs. In corpora, some verbs more likely to appear non-adjacent to their particle than other verbs. As the human comprehension system has been hypothesized to be *tuned* to the frequencies of its environment (Mitchell et al., 1996), these corpus preferences could influence reading times. We therefore predict these frequencies affect reading times. Unfortunately, it is difficult to calculate these preferences as many of the verbs used in the experiment too infrequently in even a large corpus such as the 100 million word British National Corpus (BNC). Therefore, we estimated this tendency using counts derived from the world wide web, which have proven to be a reliable method for overcoming sparse bigram statistics (Keller et al., 2002). Keller et al. (2002) observe that web counts *decrease* data sparseness at the expense of *increased* variance. Google searches were used to determine how often the word pairs were seen together or apart. For example, the pair of queries (both in quotes), "prop up" and "prop * up" were submitted, resulting in an estimate of 1.59 million and 301,000 documents. We then calculate

$$P_{together} = \frac{1.59}{1.59 + 0.301} = .841$$

A limitation of Google counts is that they do not take word class into consideration. For example, a search may well return pages with the target words which nonetheless are not in a particle verb construction, for example a sentence such as 'he won victory over himself'[2]. It is therefore necessary to verify that the Google counts are reasonable estimates for actual verb particle frequencies. To determine if this was they case, we used a version of the BNC which was parsed by a highly accurate automatic parser (Charniak, 2000). This parsed BNC was then used to derive accurate word counts for the subset of particle verbs which did appear in the BNC. Computing the Pearson correlation of the log of these BNC counts against the log of the Google counts, we found a significant Pearson correlation for the numerators, $r = .71, n = 13, P < 0.01$, and the correct trend for the denominators $r = .42, n = 13, P = 0.15$

The probability $P_{together}$ is calculated for both verbs in each item. These probabilities are referred to as vp1pref and vp2pref for the first and second verb, respectively. Including these two quantities in the regression, we get the following equation:

$$\begin{aligned} r_4 \quad \sim \quad & \mathsf{syl1} * \mathsf{syl2} * \mathsf{vp1pref} * \mathsf{vp2pref} \\ & + \mathsf{absdiff} * \mathsf{vp1pref} * \mathsf{vp2pref} \end{aligned}$$

[2] This example was provided by an anonymous reviewer.

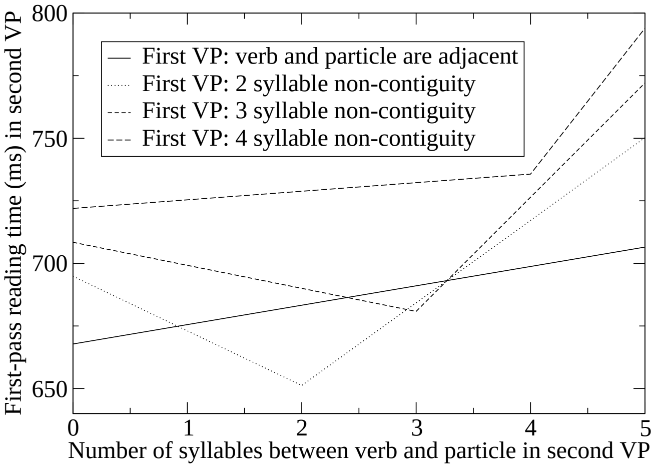

Figure 1: The predicted effect of syl2 on VP_2 first-pass reading time, when vp1pref and vp2pref take their median values. The lines show various values of syl1. There is a local minimum/inflection point at syl2 $= 0$ and when syl1 $=$ syl2.

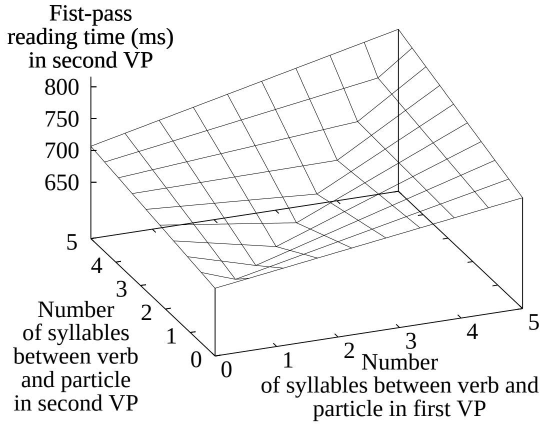

Figure 2: The predicted joint effect of syl1 and syl2 on VP_2 first-pass reading times, when vp1pref and vp2pref take their median values.

1	1st person singular pronoun
2	2nd person singular pronoun
3	3rd person singular pronoun
4	Proper names
5	Kin terms
6	Animate beings
7	Concrete objects
8	Containers
9	Locations
10	Sensual Entities
11	Abstract Entities

Table 2: Silverstein Hierarchy (SH), shown by Gries (1999) to influence particle placement

With the verb-particle preferences included, the regression now yields statistically significant estimates for nearly all coefficients in first-pass reading times. Of particular interest, the syllabic interactions syl1:syl2:vp1pref:vp2pref and absdiff:vp1pref:vp2pref were both statistically significant ($P < 0.05$). A summary of the co-efficients is shown in Table 1. No other reading time measure produced significant results.

The regression predicts ease of processing in a manner consistent with the phonological parallelism effect. While difficult to assess directly, it is possible to interpret the regression graphically. Figure 1 shows a two-dimensional characterization of the results. Choosing the median values for vp1pref and vp2pref, this figure shows reading times as a function of syl2. Each line represents a different value for syl1. As expected, there is a minimum when syl2 $= 0$, when the verb and particle are adjacent. Reading times increase as the number of syllables increase, validating earlier research showing long NPs between a verb and particle are harder to process. However, this increase in processing load is mitigated at a local minimum/inflection point when syl1 $=$ syl2, denoting syllabic parallelism. This interaction can be seen more clearly in Figure 2, which plots the relationship between syl1, syl2 and reading times in a three dimensional plot. As in Figure 1, vp1pref and vp2pref are set to their median values.

Alternative Encoding There is another explanation for the results. For both syl1 and syl2 (let us refer to them collectively as syl), the difference between keeping the verb and particle adjacent (i.e. syl $= 0$) and non-adjacent (i.e. syl > 0) is more noticeable than adding one syllable to the object noun phrase in the non-adjacent condition (e.g. changing syl $= 3$ to syl $= 4$). Therefore, it is possible that a categorical variable encoding adjacency vs non-adjacency can explain all the variance, and the ordinal variables syl and absdiff are not required. If we encode these variables as vp1 and vp2, it is possible to regress these against reading times to test this hypothesis. However, the regression

$$r_4 \sim \mathsf{vp1} * \mathsf{vp2} * \mathsf{vp1pref} * \mathsf{vp2pref}$$

yielded no significant effects for any coefficient in any reading time measure. In particular all the interaction terms vp1:vp2 and vp1:vp2:vp1pref:vp2pref were not significant ($P > .1$ in all measures).

Removing the vp1pref and vp2pref variables, it is possible to perform a traditional ANOVA analysis on the data. This also produced no statistically significant results ($P > .05$ for all measures).

Discussion

Our main result is that the reading times of particle verbs do exhibit the parallelism effect. Surprisingly, we found no evidence for pure syntactic parallelism. In other words, there was no evidence that one use of the rule VP→V NP Part predicted another use of VP→V NP Part. Rather, the parallelism effect here appears to operate on the phonological level: we found evidence for more parallelism when the number of syllables in the first NP were closer to the number of syllables in the second.

	Estimate	t value	$Pr(>\|t\|)$	
(Intercept)	-11850	-2.396	0.0168	*
syl1	-13097	-1.697	0.0901	.
syl2	-15831	-2.008	0.0449	*
vp1pref	12789	2.534	0.0115	*
vp2pref	13469	2.542	0.0112	*
absdiff	14981	1.960	0.0503	.
syl1:syl2	9745	2.339	0.0196	*
syl1:vp1pref	13468	1.703	0.0890	.
syl2:vp1pref	16380	2.027	0.0430	*
syl1:vp2pref	15084	1.822	0.0688	.
syl2:vp2pref	18079	2.137	0.0329	*
vp1pref:vp2pref	-13777	-2.549	0.0110	*
vp1pref:absdiff	-15511	-1.980	0.0481	*
vp2pref:absdiff	-17171	-2.095	0.0364	*
syl1:syl2:vp1pref	-10053	-2.352	0.0189	*
syl1:syl2:vp2pref	-11099	-2.466	0.0139	*
syl1:vp1pref:vp2pref	-15515	-1.828	0.0679	.
syl2:vp1pref:vp2pref	-18705	-2.156	0.0313	*
vp1pref:vp2pref:absdiff	17785	2.117	0.0345	*
syl1:syl2:vp1pref:vp2pref	11453	2.480	0.0133	*

Table 1: Regression coefficients and their significance. Significance codes: '*': 0.05 '.': 0.1

Moreover, it appears as if the parallelism effect is influenced by the tendencies of particle verbs to appear adjacent or not adjacent to their particles. It is possible that these preferences are due to some latent factors which have no other effects. However, another explanation is that these adjacency preferences are due to the selectional preference the verbs place upon their NP objects. In turn, characteristics of the selected NPs influence particle placement. In particular, it is known that the preference for a non-adjacent placement of a particle is influenced by the accessibility of the non-adjacent construction, which is in turn influenced by what is known as the *entrenchment* of the noun phrase (Gries, 1999). This is encoded in the Silverstein hierarchy (SH), which posits a scale of entrenchment from abstract entities down to the 1st person singular pronoun. Following Gries (1999), we utilize Deane's (1987) modified version of the SH, shown in Table 2.

Using only corpus studies, it would be tedious to directly ascertain how a verb's selectional preferences (as encoded using the SH) influence its preference for preferring an adjacent particle. However, it is possible to indirectly measure this – by encoding each object noun phrase in VP_2 with its level on the SH, and correlating this with vp2pref. We found a statistically significant Pearson correlation ($r = .53, n = 28, P < 0.01$). Although this indicates some interaction is present, it is difficult to draw firm conclusions about the nature of the interactions because this is only an indirect measure. However, it is suggestive of a semantic influence on parallelism, in addition to the phonological effect shown above.

Theoretical Implications

There has been some debate concerning the mechanism responsible for the parallelism effect. Some (Dubey et al., 2005, 2006) have argued that the parallelism effect is an instance of comprehension-comprehension syntactic priming. Others (Frazier et al., 2000) have argued for a distinct mechanism, separate from syntactic priming. Frazier et al. (2000) base their argument upon an experiment showing a null effect for priming in non-coordinate contexts whereas similar items do in fact show a parallelism effect in coordinated contexts. Dubey et al. (2005) counter this with corpus studies showing that, at least in corpora, effects can be seen both in coordinated and non-coordinated context, although the effect is weaker in the latter. As this experiment found the presence of a parallelism effect only after taking phonology and latent preference into account, a full test of a similar effect in a non-coordinated context may also depend upon taking these factors into account, in order to overcome the much smaller effect size in such contexts.

Regardless of the cause of the parallelism effect, this study does have implications upon whatever mechanism is responsible for it: parallelism does not operate on the level of syntax alone (as depicted in Figure 3(a)). In fact, Figure 3(b) gives a more accurate picture: the parallelism effect must at least operate on both the level of syntax and phonology. Moreover, there must be some interaction between the two: a change in syl1 or syl2 from 0 to a non-zero value indicates a syntactic difference, whereas other changes in syl1 and syl2 are due to phonological differences. Therefore, the parallelism effect is not purely syntactic, but operates on several distinct modules of the human sentence processor, and requires interactions between these modules.

References

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, *18*, 355–387.

Branigan, H. P., Pickering, M. J., Liversedge, S. P., Stewart, A. J., & Urbach, T. P. (1995). Syntactic priming - inves-

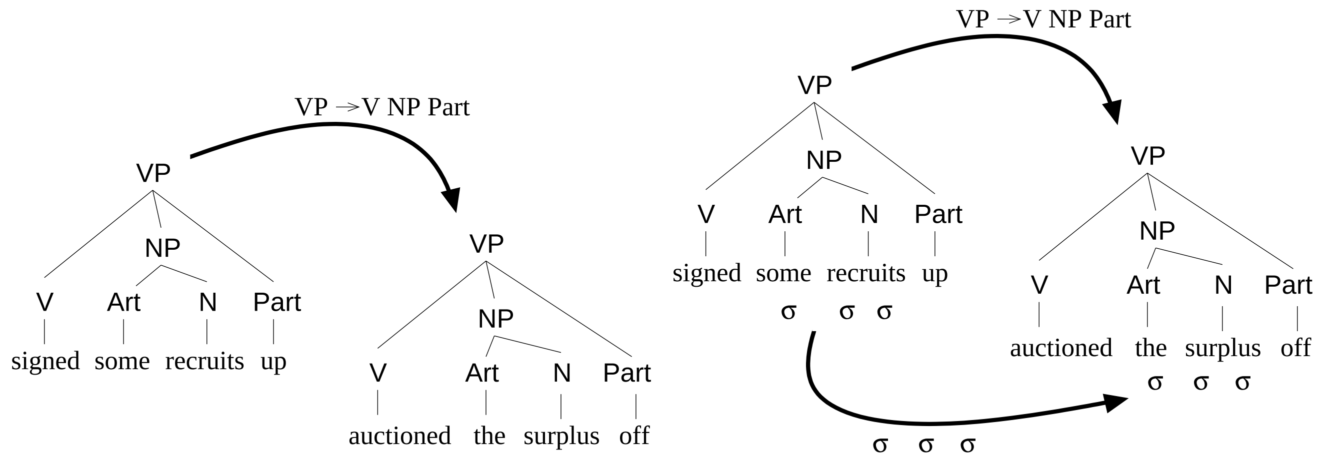

(a) The Parallelism Effect has previously been associated with syntactic structures.

(b) The present work finds evidence that the Parallelism Effect operates *simultaneously* on syntax and phonology.

Figure 3: Parallelism may operate on more than one level.

tigating the mental representation of languge. *Journal of Psycholinguistic Research*, *24*, 489–506.

Branigan, H. P., Pickering, M. J., & McLean, J. F. (2005). Priming prepositional-phrase attachment during comprehension. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *31*, 468–481.

Charniak, E. (2000). A Maximum-Entropy-Inspired Parser. In *Proceedings of the 1st Conference of North American Chapter of the Association for Computational Linguistics*, (pp. 132–139), Seattle, WA.

Deane, P. D. (1987). English possessives, topicality, and the SH. In *Proceedings of the 13th Annual Meeting of the Berkeley Linguistics Society*, (pp. 64–76), Berkeley, CA. Berkeley Linguistics Society.

Dubey, A., Keller, F., & Sturt, P. (2006). Integrating syntactic priming into an incremental probabilisti parser, with an application to psycholinguistic modeling. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, (pp. 417–424), Sydney. The Association for Computational Linguistics.

Dubey, A., Sturt, P., & Keller, F. (2005). Parallelism in coordination as an instance of syntactic priming: Evidence from corpus-based modeling. In *Proceedings of the Human Language Technology Conference and the Conference on Empirical Methods in Natural Language Processing*, (pp. 827–834), Vancouver. The Association for Computational Linguistics.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, *6*, 291–325.

Frazier, L., Munn, A., & Clifton, C. (2000). Processing coordinate structure. *Journal of Psycholinguistic Research*, *29*, 343–368.

Gibson, E. A. F. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, *68*, 1–76.

Gries, S. T. (1999). Particle movement: A cognitive and functional approach. *Cognitive Linguistics*, *10*, 105–145.

Gries, S. T. (2005). Syntactic priming: A corpus-based approach. *Journal of Psycholinguistic Research*, *34*, 365–399.

Hawkins, J. A. (1995). *A performance theory of order and constituency*. Cambridge, UK: Cambridge University Press.

Keller, F., Lapata, M., & Uryupyna, O. (2002). Using the web to obtain frequencies for unseen bigrams. In *Proceedings of EMNLP 2002*, Philadelphia, PA.

Mitchell, D. C., Cuetos, F., Corley, M. M. B., & Brysbaert, M. (1996). Exposure-based models of human parsing: Evidence for the use of coarse-grained (non-lexical) statistical records. *Journal of Psycholinguistic Research*, *24*, 469–488.

Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language*, *39*, 633–651.

Reitter, D., Moore, J., & Keller, F. (2006). Priming of syntactic rules in task-oriented dialogue and spontaneous conversation. In R. Sun (ed.), *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, (pp. 685–690), Vancouver.

Szmrecsanyi, B. (2005). Creatures of habit: A corpus-linguistic analysis of persistence in spoken English. *Corpus Linguistics and Linguistic Theory*, *1*, 113–149.