Alignment in Syntactic Blending[*]

Elizabeth Coppock
Stanford University

Using a large corpus of spontaneously uttered syntactic blends (e.g. *cast into question* from the sources *call into question* and *cast into doubt*), two main hypotheses regarding blend formation are compared: the *sequencing hypothesis*, according to which blends maximize the transition probabilities between the words or parts of speech in the sources, and the *alignment hypothesis*, which posits an abstract process of alignment as the basis for blend formation. It is found that randomly-generated unobserved blends can be reliably distinguished from observed blends on the basis of the degree to which they violate the predictions of the alignment hypothesis, although transition probabilities also play a role.

1. Introduction

When planning an utterance, people consider multiple alternative ways of expressing themselves, and these alternatives are considered in parallel. The fact that one does not discard one alternative plan before considering another is shown by speech errors like (1) and (2), in which alternative formulations of the same message are intertwined in speech, showing that they are simultaneously under consideration (examples from Garrett 1980). This type of speech error seems to involve not one but two underlying targets; these are shown in (a) and (b).

(1) sudden quicks ... stops aren't so bad.
  a. sudden stops
  b. quick stops

(2) Ed Borkowski's not letting the pressure on here.
  a. ... not letting the pressure up
  b. ... keeping the pressure on

As Garrett (1980) described it, "the message being communicated is constant but its expression is divergent;" in other words, a single meaning is being expressed

[*] Thanks to my advisor Dan Jurafsky and the participants in the workshop.

MIT Working Papers in Linguistics 53, 123–456
The State of the Art in Speech Error Research: Proceedings of the LSA Institute Workshop
Carson T. Schütze and Victor S. Ferreira (Eds.)

in two simultaneous and conflicting ways. We will call these examples *syntactic blends*. Alternative plan errors like syntactic blends are a window into the parallelism of the speech production system, because they reveal alternatives that are simultaneously present. The very existence of blends shows that syntactically and lexically diverse alternatives can be simultaneously available to the speech production system. The form that blends take is also revealing: as this paper will argue, syntactic blends are formed on the basis of correspondences between the components of alternative plans.

Researchers of competing plan phenomena face special methodological difficulties. The primary challenge is identifying the "interacting units," in Dell's (1986) terminology. In errors like word exchanges as in (3), the interacting units are the clearly the exchanged words (*rule* and *word*).

(3) words of rule formation
(intended: *rules of word formation*)

In blends, the interacting units are not located within one target, and there are multiple pairs of interacting units. In (2), for example, there appears to be an interaction of some sort between *on* and *up*, as well as an interaction between *keeping* and *letting*.

Another challenge for researchers in this domain is that there have not been enough examples of syntactic blends recorded in a suitable format for analysis. The only published analysis of blend structure uses a corpus of 64 examples (Fay 1982). A much larger corpus of syntactic blends with nearly 2000 examples was published in 1987 by Gerald Cohen in his collection entitled *Syntactic Blends in English parole*, which was converted into an electronic format for the purposes of the analyses presented here. Cohen's corpus consists of 1993 examples collected in natural speech by the author, each annotated by the author with putative targets, along with occasional contextual and speaker information, and various other observations.

To illustrate the flavor of this text, consider example #409 from Cohen's corpus in Figure 1. Cohen lists the blend itself (which I will refer to as the blend "name"), followed by the utterance in which it was heard (the "utterance"). For each example, Cohen also lists the putative sources, a and b (and sometimes but very rarely c).

```
409.  'carry through easily'   in:  'The nomination will carry
      through easily.'
           a.   The nomination will carry easily.
           b.   The nomination will go through easily.
       (And possible influence:  carry through in e.g. They carried
      the supplies through the jungle.
```

Figure 1
Example #409 from Cohen's (1987) corpus

Using Cohen's (1987) rich collection of syntactic blends and a new method for identifying the interacting units involved in their formation, this paper supports the hypothesis that syntactic similarity plays a role in blend formation.

2. Models of blend formation

Syntactic blends are restricted in form. Almost all blends can be described either as *splices* or *substitutions*, and most can be described as both (Fay 1982). A substitution blend is formed by the substitution of a single word in one target for a single word in another target. A splice blend is formed by taking an initial sequence from one target and a final sequence from another target. In addition to these two basic types, Fay also defined two types of blends derived from the basic types: *indeterminate* blends and *complex* blends. Indeterminate blends are those which can be analyzed either as splice blends or as substitution blends. For example, consider (4). It could be analyzed either as the substitution of *accept* for *believe*, or a concatenation of the initial substring up to the vertical bar in of source (a) onto the final substring beginning at the vertical bar of source (b). Similarly, in (5), we have two ways of analyzing the example, either as a substitution involving the bolded words, or as a splice of *buy* with *on clothes*.

(4) people who **accept** | in the Gospel
- a. people who **accept** | the Gospel
- b. people who **believe** | in the Gospel

(5) a little less money to **buy** | on clothes
- a. a little less money to **spend** | on clothes
- b. a little less money to **buy** | clothes

Complex blends are the opposite of indeterminate blends: they cannot be analyzed either as splice blends or substitution blends. Figure 2 shows the percentages of the various blend types in the two corpora. Only around 5% of each corpus is constituted by "complex" blends; "indeterminate" blends are the most common type by far, and substitution and splice blends are somewhere in the middle.

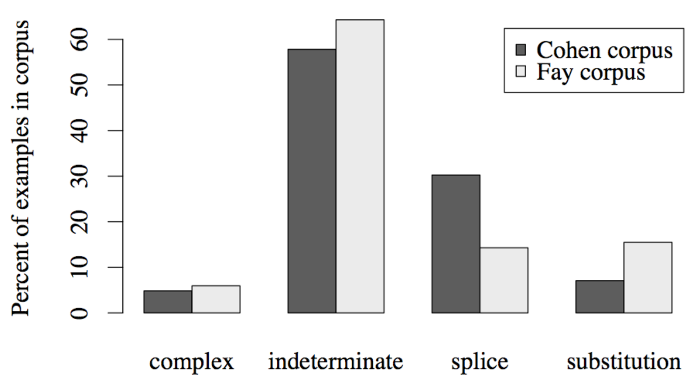

Figure 2
Distribution of blend types in Fay and Cohen corpora

Fay (1982) also shows that in all of the unambiguous substitution blends, the substituted word is always of the same part of speech category as the word it substitutes for. Thus blend formation does not arise from an arbitrary mish-mash of the two sources; rather, there are regularities underlying the formation of blends.

2.1 Fay's prosodic alignment hypothesis

For splice blends, Fay argues that the choice of splice point (hence, to some degree, the alignment) is influenced by prosody "rather than" syntax on the basis of the following two examples:

(6) It doesn't make any matter.
  a. It doesn't make any difference.
  b. It doesn't matter.

(7) if you're not watch out
  a. if you're not careful
  b. if you don't watch out

In (6), the interacting units are *matter* and *difference*; in (7), *careful* and *watch out*. These are not syntactically, but prosodically parallel units. Note that *watch out* and *careful* are not in drastically different syntactic positions; they are both predicates governed by an auxiliary and in the scope of negation, so (7) doesn't make a terribly powerful argument against the role of syntax in alignment. (6) presents a stronger case; a syntactic alignment would slot *matter* with *make*, not *matter* with *difference*. The prosody of *matter* and *difference* is (presumably) similar, though, because they both (presumably) bear nuclear accent.

Fay gives a model of splice-formation which includes an algorithm for aligning the sources prosodically. The algorithm goes as follows. First, mark the words in each target that bear primary and secondary stress. Then align the sentences so that primary stressed words occur in the same position and likewise for secondary stress. For example, the alignment between *What the hell am I doing here?* and *What the hell am I here for?* (which happened to produce *What the hell am I doing here for?*) would be as follows:

(8)

Whàt	the	hèll	am I	dóing	here?
Whàt	the	hèll	am I	hére	for?

Given this alignment two types of blends may occur. A *pre-stress* splice occurs immediately before the following stressed word (e.g. *It doesn't make any | matter*) In a *post-stress* splice, the following unstressed words from both targets are incorporated into the error (e.g. *What the hell am I doing here | for?*).

2.2 The syntactic alignment hypothesis

Fay's proposal is a type of *alignment and selection* model; it includes a process of *alignment*, which creates a sequence of slots, where each slot corresponds to a lexical item (or perhaps a short sequence of words), and a *selection* process, which chooses exactly one element from each slot. Within this class of models, there could be several hypotheses.

The goal of the present paper is to support a *syntactic*, rather than prosodic, alignment hypothesis, according to which syntactically parallel units align. Take the following example:

(9) They should lend a little hand.
  a. They should lend a little help.
  b. They should give a hand.

Here, *lend* is in a parallel syntactic position to *give*, as the main verb of the sentence, and *hand* and *help* are in parallel syntactic positions as the head noun of the direct object of the main verb. If words in syntactically similar positions align with one another, then these pairs of words (*lend* and *give*, and *hand* and *help*) should align with one another. The resulting blend should therefore be one of: *They should lend a little hand* (which we get), *They should give a little hand*, *They should give a little help* (which wouldn't be classified as a blend, since it's identical to one of the sources), and perhaps *They should lend a help* (if *a* and *a little* align, as determiner phrases). We should not get *They should lend a hand help*, or *They should give lend a hand*, for example. In fact, example (9) could be accounted for without any reference to headedness, with only a constraint that words of the same part of speech align. Call this constraint the *Weak Syntactic Alignment Hypothesis*: words of the same part of speech align. This is the hypothesis for which the present paper gives evidence.

There is already some support for the syntactic alignment hypothesis.[1] Recall that Fay (1982) shows that in all of the (unambiguously) substitution blends, the substituted word is always of the same part of speech category as the word it substitutes for. Thus, Fay shows that, to a limited extent, syntactic category plays a role in blend structure. In general, the interacting units in word exchanges tend to be of the same part of speech category (Garrett 1975).

Cutting and Bock (1997) found that both semantic and syntactic similarity facilitate blending. In their study, speakers were asked to produce one of the two idioms they had recently been presented with, in three conditions: (i) the two idioms had the same syntax and the same meaning (e.g. "shoot the breeze" and "chew the fat"), (ii) the same syntax and a different meaning (e.g. "shoot the breeze" and "raise the roof"), or (iii) different meaning and different syntax (e.g. "shoot the breeze" and "nip and tuck"). Speakers produced idiom blends in 1.7% of the cases overall (in addition to other types of errors). It was found that speakers produced significantly more blends in the same-meaning/same-syntax condition than in the different-meaning/same-syntax condition, and more blend errors in turn in the different-meaning/same-syntax condition than in the different-meaning/different-syntax condition. This gives evidence that both syntactic and semantic similarity influence the likelihood that two sources will blend.

[1] Fay (1982) argues that syntax does not play a role in determining where the splice will occur in "splice" blends. To argue that syntax does not play a role in determining the splice point, Fay (1982) states that "In only 8 of the 19 cases [of splice blends] did a perfect structural parallel [between the sources] exist" (p. 728). The lack of a perfect structural parallel between the sources merely makes it harder to specify the predictions of a syntactic alignment hypothesis, but this does not provide evidence against such a hypothesis.

One would expect syntactic and semantic similarity to play a role in the determination of blend form in light of these results, because the correspondences between elements of the alternative plans that contribute to overall syntactic similarity give structural relationships that may characterize the form that blends take. Cutler (1982a) points out that phonological similarity in Freudian slips not only facilitates them, but also helps determine the formal shape of the error. If syntactic similarity in blending functions similarly to how phonological similarity functions in Freudian slips, we should expect it not only to facilitate blending, but also to determine the form that it takes.

To account for category-preservation effects like this, models of speech production often introduce a distinction between frames containing slots and the content that fills them. A category-preservation effect that has received slightly more attention is the fact that in phonological errors, onsets, for example, tend to interact with onsets; slot/filler models of this effect make use of a phonological frame with slots for the onset, the nucleus and the coda of the current syllable, and a corresponding categorization of the potential fillers (Dell 1986, 1988; Levelt 1989; MacKay 1987; Meyer 1990, 1991; Reich 1977; Shattuck-Hufnagel 1979, 1987; Stemberger 1990, 1991 i.a.).[2] In the case of syntactic category-preservation effects the division is between syntactic frames consisting of slots with specified categories on one hand, and lexical items to fill the slots on the other hand (Bock 1982, 1986; Dell 1986; Garrett 1975, 1980; Kempen and Hoenkamp 1987; Levelt 1989). With a mechanism that prevents fillers from being inserted into a slot of the wrong type, this separation accounts for the category-preservation effect in word exchanges and substitutions. (Models without an explicit distinction between slots and fillers, such as the PDP models described by Dell et al. (1993), are capable of modelling category-maintenance effects and are also consistent with rare occasional violations of them.)

For instance, in the Dell (1986) model, when a slot in the syntax with a given part of speech (e.g. *verb*) is ready to be filled, the most active member of that category is selected. This model posits syntactic frames that contain indexed slots at the terminal nodes. The frame slots are co-indexed with word nodes in the connectionist network via indices, and the most active node fills the slot. One possible way to extend this model for the case when multiple plans are active would be be to allow multiple frames, each with their own set of slots. Word nodes may be associated with one slot in one frame, but another slot in another plan, with multiple flags, one for each plan the node is associated with. Regardless of how it is done, any generalization of this model to account for multiple plans should predict interactions between elements of the same lexical category, even across plans. In fact, the only type of interaction that the Dell (1986) model, modified in such a way, would allow is between words of the same category.

[2] In a slightly different approach to the problem of "binding" a filler to a slot, Roelofs (1997) proposes a "checking" mechanism attached to filler nodes that ensures that the filler is in the right slot. Speech errors arise when this checking mechanism fails to be careful enough, ensuring only that the category of the slot is appropriate without verifying the identity of the slot. Roelofs' solution is one of *binding-by-checking*, as opposed to *binding-by-timing*, as in Dell's (1986) model. These two approaches both predict interactions between same-category units.

2.3 Sequencing

An alternative hypothesis regarding blend formation is that the resulting string maximizes the transition likelihoods between either the words in the sources (Word Sequencing Hypothesis), or the parts of speech of the words in the sources (POS Sequencing Hypothesis). According to this hypothesis, if pairs of verbs, for example, tend to be uniquely realized, this is not due to an alignment between the verbs, but rather to the fact that verb−verb sequences are highly unlikely.

3. Logistic regression study

To test the syntactic alignment hypothesis, a set of pseudo-blends was automatically generated by merging together arbitrary substrings of the sources, and a logistic regression study was carried out in which the dependent variable was whether or not the candidate blend was actually observed or not. Alignment and sequencing constraints were evaluated for each blend. If alignment plays a role in blend formation, then the number of alignment constraints violated by a candidate blend make it less likely to be one that was observed.

3.1 Procedure

The majority of the Cohen text was scanned in, Optical Character Recognition (OCR) was performed on the scanned images, and the mistakes were corrected partially automatically. To complete the resulting database, I identified the subpart of the sources corresponding to the blend "name" by hand. I also hand POS-tagged each of the reduced sources using the Penn Treebank guidelines.

Several types of examples in the Cohen corpus were excluded from analysis. Those spoken by children or non-native speakers were not included because they may reflect a non-standard underlying grammar. Those involving segmentations within a word boundary (e.g. *when every time I make a speech*, from *whenever I make a speech* and *every time I make a speech*) were also excluded due to the difficulty of analyzing them (not because they seem qualitatively different, necessarily; this is an issue for further exploration). Examples that sounded fully acceptable to me were also excluded, but I did not quibble extensively with Cohen's judgments, in order to avoid bias. Finally, for the sake of computation time, I only included blends whose sources contained fewer than five words.

For each remaining blend, a set of pseudo-blends was generated. This set contained every possible merge of every single substring of the two sources, allowing deletions. The only constraints on the pseudo-blends were that (1) they preserve the order of the words in the sources and (2) they contain at least one word from each of the sources. Each pseudo-blend received four scores:

- An n-gram language model score estimating the transition probabilities between the **words** in the blend or pseudo-blend,
- An n-gram language model score for the **parts of speech** in the example,
- The number of **lexical** alignment constraints the candidate violates,
- The number of **POS** alignment constraints violated by the candidate.

These scores were used as predictors in a logistic regression model whose dependent variable was whether or not the blend was actually observed. Because the number of pseudo-blends was so enormous, one pseudo-blend per observed blend was randomly chosen for analysis, to balance the sizes of the observed and non-observed classes.

The alignment constraints were derived using a program called SCLITE, which is normally used for scoring the outputs of speech recognizers. SCLITE provides an alignment of two strings which minimizes the number of insertions, deletions, and substitutions, maximizing the number of common elements. This program was used to align the sources (corresponding to the blend "name") with each other. The POS-tag sequences corresponding to the sources were also aligned with one another. The common elements in the alignment between two sources for a blend correspond to alignment constraints.

To test whether a given alignment constraint was violated or not, correspondences between the source strings and the candidate blends were identified, again using SCLITE. Common elements in the alignment between a source and the blend were considered corresponding elements. This information was supplemented by a trace of the program that generated the blend, which gave additional correspondences between the sources and the blend. These correspondences were used in the algorithm described Appendix A..

3.2 Results and discussion

All four of the predictors (lexical n-gram language model score, POS n-gram language model score, lexical alignment, and POS alignment) made a highly significant improvement in the model's ability to distinguish observed blends from pseudo-blends, in each of 20 trials in which one pseudo-blend was randomly chosen for each observed blend. The full logistic regression model for predicting the log of the odds that a blend was observed had the following coefficients on average:

$$\text{Prob}\{\text{observed} = 1\} = \frac{1}{1 + \exp(-X\beta)}, \ \text{ where}$$

$$\begin{aligned} X\hat{\beta} &= 4.272 + 0.008 * \text{ngram\_lex} + 0.0307 * \text{ngram\_pos} \\ &\quad -0.754 * \text{align\_lex} - 1.929 * \text{align\_pos} \end{aligned}$$

In this equation:

- `observed` is the dependent variable, whether or not the blend was actually observed,
- `ngram_lex` refers to the language model score (a log probability) given to the the sequence of words in the candidate blend by the lexical n-gram language model,
- `ngram_pos` refers to the language model score given to the sequence of parts of speech of the candidate blend given by the POS n-gram language model,
- `align_lex` is an integer ranging from 0 to 4 representing the number of lexical alignment constraints violated by the candidate blend, and

- `align_pos` is the number of POS alignment constraints violated by the candidate blend.

The mean classification accuracy of this model over 20 trials is 83.31%, with a standard deviation of 0.48%, where classification accuracy is defined as the percent of candidates assigned a positive log odds of being observed when observed or a negative log odds of being observed when not observed. When POS alignment is removed from the model as a factor, mean classification accuracy drops to 76.06%, with a standard deviation of 0.55%. Thus, adding POS alignment improves classification accuracy of the model by 7.25%. This means that the more POS alignment constraints a candidate violates, the less likely it is to have been observed as a blend. According to the average model, the odds decrease by roughly twice the number of alignment constraints violated by the candidate (the coefficient for `align_pos` is -1.929); this far exceeds the impact of lexical alignment violations, which incur a smaller decrease in log odds for each violation (the coefficient for `align_lex` is -0.754). The graph in Figure 3 shows the average classification accuracy of the logistic regression models, with and without the POS alignment predictor variable. The difference in accuracy is significant; the 95% confidence intervals are too small to be shown on the graph.

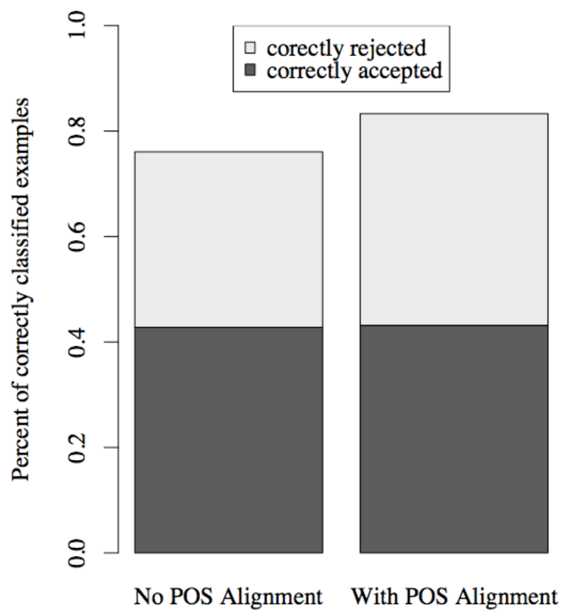

Figure 3
Classification accuracy of average regression model

As Figure 3 shows, most of the improvement in accuracy comes from correct rejections, rather than correct acceptances. This is due to an asymmetry in the data stemming from the fact that identical words usually have the same part of speech. The examples that are correctly rejected by the full model but not by the model without POS alignment violations tend to be cases in which there are POS alignment constraints on distinct lexical items. For example, a pseudo-blend generated from the sources *think about* and *consider* was *think consider about* (the observed blend was *consider about*):

(10) consider about
  a. think_{VB} about_{IN}
  b. $\text{consider}_{\text{VB}}$

The sources have part of speech tag sequences VB-IN and VB, respectively (VB stands for "verb", IN for "preposition"). The pseudo-blend *think consider about* violates the POS alignment constraint that *think* should correspond to *consider*, whereas the observed example, *consider about*, does not. The full model correctly rejects *think consider about*, but the model without POS alignment incorrectly accepts it.

There were very few examples that were correctly accepted by the full model but incorrectly rejected by the model without syntactic alignment. One of these was *by and away* (from *by far* and *far and away*), which violates the lexical alignment constraint that *far* corresponds to *far*. This example violates no POS alignment constraints, because there are none to violate; *far* is tagged as an adjective in *by far* but as an adverb in *far and away*. Another example is *for the (incredible) price for (just $41.95)*, from (A) *for the price of* and (B) *for*; this one violates the lexical alignment constraint that *for* corresponds to *for*, but no syntactic alignment constraints. Although source A's *for* has the same part of speech as source B's *for*, the part of speech tag of source B's *for* happened to be aligned with the part of speech tag of source A's *of* due to an arbitrary preference for rightward alignments in SCLITE, and this was treated as the only syntactic alignment constraint.[3] Another example of this type is *problems of these nature*, formed from *these problems* and *problems of this nature*. Here, the lexical alignment constraint that *problems* corresponds to *problems* is violated, but this is not a POS alignment constraint because the part of speech of *these* was aligned with the part of speech of *this*. But usually, lexical alignment constraints correspond to POS alignment constraints, so there are few cases in which the number of POS alignment constraints is lower than the number of lexical alignment constraints.

4. Prosodic vs. syntactic alignment

The result in the previous section supports the syntactic alignment hypothesis, but does not rule out a prosodic alignment hypothesis, since prosodic alignment constraints were not included in the model. It is not feasible in general to assign prominence levels to unspoken plans, so testing this hypothesis cannot be done automatically, using the entire corpus. However, the syntactic alignment hypothesis can be shown to make superior predictions to the prosodic hypothesis, within the domain of blends between single verbs and verb/particle combinations. In general, nuclear accent goes on the particle in the verb/particle combination, but on

[3] The fact that there can in principle be many legal alignments between two given strings is not taken into account here, and is probably a source of error.

a verb when it is phrase-final. Thus, the Prosodic Alignment Hypothesis predicts that in a blend of a single verb with a sequence containing post-verbal stress (e.g. *caught* plus *found out*), the single verb should align with the post-verbal stress (yielding *found caught*). The prediction of the Syntactic Alignment Hypothesis is that the single verb should align with the verb (yielding *caught out*).

The set of such blends is in Table 1. All of them involve verb/verb alignment; in none of them does the verb align with the particle.

Table 1
Blends formed from a single verb and a verb-particle combination

Example	*Source A*	*Source B*
If we were **caught out** ...	caught	found out
I **checked base** with David about...	checked	touched base
It **consumes up** a lot of time on the clock.	consumes	uses up
If this **continues on** ...	continues	goes on
Why don't you **choose** something **out**?	choose	pick out
I can't **fathom out** how it could have happened.	fathom	figure out
They might **fold under**.	fold	go under
Let's **list down** all the options.	list	write down
It doesn't **meet up** to his standards.	meet	measure up
They have **proved out** very successful for me.	proved	turned out
I hate to see it **returned back** to them.	returned	turned back
Don't you **start** rumors **around** like that.	start	spread around
He **slipped down**.	slipped	fell down
There are people **waiting by** to take your phone call.	waiting	standing by

We make similar predictions for blends between phrase-final intransitive verbs and complement-taking verbs. A phrase-final intransitive verb will take a nuclear accent, whereas the complement of a complement-taking verb will take nuclear stress, so the intransitive verb should align with the complement according to the prosodic alignment hypothesis, but with the verb according to the syntactic alignment hypothesis. Only the predictions of the syntactic alignment hypothesis are borne out, as shown in Table 2.

Table 2
Blends of a complement-taking verb and a nuclear accent-taking verb

Example	*Source A*	*Source B*
They were **glutted to capacity**.	glutted	filled to capacity
They can't **touch a candle to** him	touch	hold a candle to
They could **decide this decision**.	decide (this)	make this decision
This **suggests** to me **the impression**	suggests to me	gives me the impression
They didn't **control in command of** that.	control	be in command of
We must **deter the stem of** brutal crime	deter	stem the growth of

(The *deter the stem of* example is not predicted by either hypothesis; the syntactic alignment hypothesis would predict *deter the growth of*, and the prosodic alignment hypothesis would predict *stem the deter of*.)

There is not enough data for statistics, but the data "suggests the impression" that verb−verb alignment is more likely than alignment between a stressed verb and a stressed element of another category.

5. Conclusion

This paper investigated the role of alignment in predicting the form that a syntactic blend will take, given its sources. A role for two types of alignment was supported: lexical alignment (identical words tend to align) and part of speech alignment (words of the same POS tend to align). Blends that actually occur in the corpus can be distinguished from arbitrary merges of substrings of the sources, based on whether or not they adhere to part of speech alignment constraints, even after the transition likelihoods between the words and the parts of speech are taken into account, in support of the Syntactic Alignment Hypothesis.

These results provide further support for the psychological reality of syntactic classes, and also suggest that syntactic processing resources are shared across plans. Syntactic alignment effects are vaguely expected under any model of speech production that contains competition between words of the same part of speech category, but more work needs to be done before this effect can be precisely derived. Posing a challenge for models of speech production, this research makes one step towards understanding the way that parallelism works when sentences are produced.

This research also identified some new methodological techniques that could be used in other areas of speech error research, or linguistic research more broadly. The technique of randomly generating candidate errors from the sources (or "targets" to use the more prevalent term), and using features to classify the errors as "observed" or "not observed" could be used in the study of other types of slips of the tongue as well. Moreover, the logic and the algorithm for detecting alignment constraint violations could be applied to other processes that combine strings, such as code-switching and what is referred to as "alignment" in the literature on conversational discourse. The author hopes that some nerdy readers will enjoy the Appendix.

A. Identifying alignment constraint violations

A.1 Logic

Consider the following blend.

(11) a bite to lunch
  a. a_{DT} bite_{NN} to_{TO} eat_{VB}
  b. a_{DT} bite_{NN} for_{IN} lunch_{NN}

Both of the sources are indefinite NPs headed by *bite*, but one has an infinitival relative modifier (*to eat*), and the other has a prepositional phrase modifier (*for* *lunch*). The POS tags shown after forward slashes in (11) reflect this difference.

There are two lexical alignment constraints on this blend: *a* must correspond with *a*, and *bite* must correspond with *bite*. We can denote these constraints using the following syntax:[4]

$$A_0 = B_0$$

$$A_1 = B_1$$

Aligning the parts of speech of the sources gives us the same two alignment constraints, so these two constraints are the only ones that apply to this example.

Both constraints are satisfied by the blend string, *a bite to lunch*, since *a* and *bite* are uniquely present as the first two words. To derive this conclusion logically, it is necessary to specify some facts that are given and an inference rule. The alignment between source A and the output string gives that word zero (*a*) of source A corresponds to word zero (*a*) of the output string:

$$A_0 = O_0$$

Similarly, the alignment between source B and the output string gives:

$$B_0 = O_0$$

With the following inference rule:

$$A_x = O_z \wedge B_y = O_z \Rightarrow A_x = B_y$$

we can derive that $A_0 = B_0$. We can derive that $A_1 = B_1$ in the same way.

Consider the pseudo-blend *to eat for lunch*. This one satisfies neither alignment constraint, as intuitively suggested by the fact that *a* and *bite* are not present in the string. Let us derive this conclusion. If A_0 is aligned with B_0, then there must be some word in the output string O_c that they both correspond to:

$$A_i = B_j \Rightarrow A_i = O_c \wedge B_j = O_c,$$

for a new constant c. Next, observe that word 2 of source A, *to*, is aligned with word 0 of the output string, *to*:

$$A_2 = O_0$$

Using this fact, we can derive the impossible conclusion that word 0 of source A (A_0) is aligned with something below zero. We can derive this conclusion using the following inference rule:

$$A_i = O_c \wedge A_j = O_k \wedge j > i \Rightarrow k > c$$

For future reference, call this the "little sister rule," which expresses that i is the "little sister" of a j that is aligned with k. Substituting 0 for i, 2 for j, and 0 for k, this gives:

$$A_0 = O_c \wedge A_2 = O_0 \wedge 2 > 0 \Rightarrow 0 > c$$

[4] If a semantics is desired, the expression $A_i = B_j$ can be interpreted to mean that the pair $\langle A_i, B_j \rangle \in \mathcal{A}$, where $\mathcal{A}$ is the set of pairs of words that are aligned.

All of the antecedents are satisfied: $A_0 = O_c$ by assumption, it is given that $A_2 = O_0$ and clearly $2 > 0$. So, the conclusion $0 > c$ follows. It's impossible for any index to be below 0 (or above the last index of the string), which we can express with the following principle:

$$\forall x \neg (0 > x)$$

In particular:

$$\neg (0 > c)$$

Thus, we have both $0 > c$ and $\neg(0 > c)$. Since $P \wedge \neg P \Rightarrow \perp$, we have derived a contradiction from assuming that $A_0 = B_0$, so this constraint is not satisfied.

We can construct a more general proof of the same conclusion, without using the fact that zero is special. This more general proof uses the given fact that the substring *a bite* of source B matches up with the substring *to eat* of the output string, which is derived from the the alignment of source B with the output string using a separate algorithm. Let us call an alignment of a substring with a substring a "free zone," expressing the idea that we know some upper and lower bounds on what aligns with what, but we don't know exactly what aligns with what. (The two substrings involved in a free zone need not be of the same length, as they are in this particular example.) Let us denote a free zone thus:

$$A_i ... A_n = O_j ... O_m$$

where $A_i ... A_n$ represents the substring of source A containing all of the A_x such that $i \leq x < n$. We are given that:

$$B_0 ... B_2 = O_0 ... O_2$$

This means that if B_0 is aligned with c, then c is in the range $O_0 ... O_2$. Stated generally:

$$A_i ... A_n = O_j ... O_m \wedge A_i = O_c \Rightarrow j \leq c < m$$

Substituting B_0 for A_i, B_2 for A_j, 0 for j and 2 for m, we have:

$$B_0 ... B_2 = O_0 ... O_2 \wedge B_0 = O_c \Rightarrow 0 \leq c < 2$$

Since the antecedents are satisfied (by assumption), we have:

$$0 \leq c$$

We can use the little sister rule in the same way that we did above to conclude that $c < 0$, and again we have a contradiction.

The general schema for proofs that an alignment constraint ($A_i = B_j$) is violated is to assume that there is some O_c to which both A_i and B_j correspond, and show either of the following:

- $c = x \wedge c = y \wedge x \neq y$
- $x \leq c \wedge c < y \wedge x \leq y$.

For the first strategy, statements of the form $c = x$ or $c = y$ can be derived through correspondence constraints that are given by the string alignment algorithm:

(a) $A_i = O_x \wedge A_i = O_c \Rightarrow c = x$
(b) $B_j = O_x \wedge B_j = O_c \Rightarrow c = x$

Statements of the form $y \leq c$ can be derived in three ways, using the assumptions that $A_i = O_c$ and $B_j = O_c$:

- Source-output constraints:
  (a) $A_i = O_c \wedge A_i = O_y \Rightarrow y = c \Rightarrow y \leq c$
  (b) $B_j = O_c \wedge B_j = O_y \Rightarrow y = c \Rightarrow y \leq c$

- Free zone rules:
  (a) $A_i = O_c \wedge A_m ... A_n = O_y ... O_x \wedge m \leq i < n \Rightarrow y \leq c$
  (b) $B_j = O_c \wedge B_m ... B_n = O_y ... O_x \wedge m \leq j < n \Rightarrow y \leq c$

- Big sister rules:
  (a) $A_i = O_c \wedge A_j = O_y \wedge j < i \Rightarrow y < c \Rightarrow y \leq c$
  (b) $B_j = O_c \wedge B_j = O_y \wedge j < i \Rightarrow y < c \Rightarrow y \leq c$

Statements of the form $c < x$ can be derived in two ways:

- Free zone rules:
  (a) $A_i = O_c \wedge A_m ... A_n = O_y ... O_x \wedge m \leq i < n \Rightarrow c < x$
  (b) $B_j = O_c \wedge B_m ... B_n = O_y ... O_x \wedge m \leq j < n \Rightarrow c < x$

- Little sister rules:
  (a) $A_i = O_c \wedge A_k = O_x \wedge i < k \Rightarrow c < x$
  (b) $B_j = O_c \wedge B_k = O_x \wedge j < k \Rightarrow c < x$

A.2 Algorithm

In the interest of computational efficiency, I developed a constraint-checking algorithm that relies on the logic described above but which does not make it explicit. The algorithm returns TRUE if the constraint is satisfied, and FALSE if not.

1. If $A_i = O_x$ and $B_j = O_y$ and $x = y$, then return TRUE.
2. If $A_i = O_x$ and $B_j = O_y$ and $x \neq y$, then return FALSE.
3. Check big and little sister rules at extrema.
   (a) If $A_k = O_0$ and $k > i$ or $B_k = O_0$ and $k > j$, return FALSE. (Little sister rules at zero.)
   (b) If $A_k = O_g$ and $k < i$ or $B_k = O_g$ and $k < j$, where g is the last index of the output string, return FALSE. (Big sister rules at end of output string.)
4. Find the greatest y such that $y \leq c$.
5. Find the lowest x such that $c < x$.
6. If $x \leq y$, return FALSE, else return TRUE.

Subroutine: finding the greatest y such that $y \leq c$.

1. Set y equal to 0.
2. (a) If $A_i = O_{y'}$ and $y' > y$, set y to y'.
   (b) If $B_j = O_{y'}$ and $y' > y$, set y to y'.
3. Freezone rules:
   (a) If there is a freezone $A_m ... A_n = O_{y'} ... O_x$ such that $m \leq i < n$, and $y' > y$, set y to y'.
   (b) If there is a freezone $B_m ... B_n = O_{y'} ... O_x$ such that $m \leq j < n$, and $y' > y$, set y to y'.
4. Big sister rules:
   (a) For each constraint $A_n : O_{y'}$, if $n < i$ and $y' > y$, set y to y'.
   (b) For each constraint $B_n : O_{y'}$, if $n < j$ and $y' > y$, set y to y'.
5. Return y.

The subroutine for finding the smallest x such that $c < x$ is very similar, except that x begins at one index past the last index of the string and decreases when smaller x's are found, Step 2 is omitted (because we are looking for a strict less-than relationship), the upper bounds of the freezones rather than the lower bounds are used, and little sister rules are used in place of the big sister rules.

References

Bock, Katherine (1982). Towards a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review* 89, 1–47.

Bock, Katherine (1986). Meaning, sound, and syntax: Lexical priming in sentence production. *Journal of of Experimental Psychology: Learning, Memory & Cognition* 12, 575–586.

Cutler, Anne (1982a). The reliability of speech error data. In Cutler 1982b.

Cutler, Anne (ed.) (1982b). *Slips of the Tongue and Language Production*. Amsterdam: Walter de Gruyter/Mouton.

Cutting, J. Cooper and Kathryn Bock (1997). That's the way the cookie bounces: Syntactic and semantic components of experimentally elicited idiom blends. *Memory and Cognition* 25, 57–71.

Dell, Gary S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review* 93, 283–321.

Dell, Gary S. (1988). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. *Journal of Memory and Language* 27, 124–142.

Dell, Gary S., Cornell Juliano, and Anita Govindjee (1993). Structure and content in language production: A theory of frame constraints in phonological speech errors. *Cognitive Science* 17, 149–195.

Fay, David (1982). Substitutions and splices: A study of sentence blends. In *Slips of the Tongue and Language Production*. In Cutler 1982b.

Garrett, Merrill F. (1975). The analysis of sentence production. In *The psychology of learning and motivation*, vol. 9, ed. G. H. Bower. Academic Press.

Garrett, Merrill F. (1980). Levels of processing in sentence production. In *Language production, Volume I: Speech and talk*, ed. Brian L. Butterworth. London: Academic Press.

Kempen, Gerard and Edward Hoenkamp (1987). An incremental procedural grammar for sentence formulation. *Cognitive Science* 11, 201–258.

Levelt, Willem J.M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.

MacKay, Donald G. (1987). *The organization of perception and action: A theory of language and other cognitive skills*. New York: Springer.

Meyer, Antje S. (1990). The phonological encoding of successive syllables. *Journal of Memory and Language* 29, 524–545.

Meyer, Antje S. (1991). The time course of phonological encoding in language production: Phonological encoding inside a syllable. *Journal of Memory and Language* 30, 69–89.

Reich, Peter A. (1977). Evidence for a stratal boundary from slips of the tongue. *Forum Linguisticum* 2, 211–217.

Roelofs, Ardi (1997). The WEAVER model of word-form encoding in speech production. *Cognition* 64(3), 107–142.

Shattuck-Hufnagel, Stephanie (1979). Speech errors as evidence for a serial-ordering mechanism in sentence production. In *Sentence processing: Psycholinguistic studies presented to Merrill Garrett*, eds. W. E. Cooper and E. C. T. Walker. Hillsdale, NJ: Erlbaum.

Shattuck-Hufnagel, Stephanie (1987). The role of word-onset consonants in speech production planning: New evidence from speech error patterns. In *Motor and sensory processes of language*, eds. E. Keller and M. Gopnik. Hillsdale, NJ: Erlbaum.

Stemberger, Joseph P. (1990). Word shape errors in language production. *Cognition* 35, 123–157.

Stemberger, Joseph P. (1991). Apparent antifrequency effects in language production: The addition bias and phonological underspecification. *Journal of Memory and Language* 30, 161–185.

Stanford University
Margaret Jacks Hall
Building 460
Stanford, CA 94305
USA

coppock@stanford.edu
http://www.stanford.edu/∼coppock