

Parsing speech corpora

Mark Johnson (Brown University)

mark_johnson@brown.edu

This talk describes recent research into parsing speech transcripts and acoustic signals (specifically, speech recognizer output). Speech parsers face several challenges in addition to those encountered in parsing text. Word identity is often only partially recoverable from the acoustic signal. The speech signal does not come neatly segmented into words or even into sentences. The presence of speech errors and disfluencies complicates the recovery of syntactic and semantic dependencies (i.e., parsing). Statistical parsers have an advantage in such settings, since they can take advantage of weak or noisy information present in the speech signal. This talk describes the kinds of statistical methods we use to parse speech, including statistical parsing models that can exploit prosodic information and can detect and to some extent correct speech errors.

Syntactic persistence in real life (spontaneous speech)

T. Florian Jaeger (Stanford University)

tiflo@stanford.edu

While speakers' tendency to reuse syntactic structures (syntactic persistence, [1]) has been demonstrated experimentally, little is known about its role in spontaneous dialogue (but see [4, 7]). We present evidence that (a) Syntactic persistence occurs in spontaneous speech; (b) Its effects persist over a long time, which argues against transient activation accounts; and, in contrast with some accounts of persistence: (c) Across-speaker syntactic persistence is significantly weakened (or non-existent) in non-task oriented dialogue. The results are based on 6,500 complement clauses (CCs) and 3,500 relative clauses (RCs) exhibiting *that*-omission from the Switchboard corpus of phone conversations:

(1) We hope [(*that*) this study contributes to the understanding of syntactic persistence in spontaneous speech]. CC example

We used Logit Generalized Linear Mixed Models to test whether and how persistence affects *that*-omission after controlling for factors known to influence *that*-omission, including predictability, weight, accessibility, and production difficulty [5, 6]. RC and CC results are almost identical. We present the CC results.

PERSISTENCE: Syntactic persistence affects speakers' choice in *that*-omission beyond the control factors and beyond individual speakers' biases (included as random effects). In contrast to experimental findings [3], the effect is mostly driven by cases in which the preceding CC (the 'prime') contained *that*. Those CCs are about twice as likely to have *that* as CCs not preceded by another CC (the baseline; $p < .001$). Both for CCs and RCs, the less frequent / less preferred form (CCs/RCs with *that*) primes more strongly.

PRIME-TARGET DISTANCE: The persistence effect decreases logarithmically with increasing distance between target and prime. Marginal effects of syntactic persistence survive over twenty sentences. This supports implicit learning accounts of syntactic persistence [2] over transient activation accounts. First results suggest that persistence effects are not cumulative, i.e., only the *closest* prime seems to affect target realization.

LEMMA-(IN)DEPENDENCE: The persistence effect is more than doubled for prime-target pairs with the same embedding verb (*hope* above), but persistence also remains significant if the embedding verb differs ($p < .02$).

SPEAKER-(IN)DEPENDENCE: We investigated whether, in spontaneous speech, syntactic persistence is observed across speakers. While within-speaker persistence effects are highly significant, we found no significant effect of across-speaker persistence ($p > .50$). This cannot be attributed to the larger distance of across-speaker primes to the target (on average 0.7 turns or 4 words): no across-speaker persistence is found for CCs for which the closest prime was across speakers ($p > .30$). Crucially, within-speaker persistence remains significant ($p < .05$) and the effect magnitude is even increased: CCs after a within-speaker prime with *that* are about 2 times more likely to contain *that* than CCs after a prime without *that*.

In conclusion, results from spontaneous speech confirm some and potentially challenge other experimental findings. Corpus-based approaches to syntactic persistence thus complement results from experiments. The extreme weakness / absence of across-speaker syntactic persistence (similar results were obtained for RCs) is puzzling given experimental results. We originally hypothesized that the difference may be due to Switchboard speakers not knowing each other and hence being less attentive to each others dialogue contributions. Preliminary results from another variation, however, do not support this hypothesis.

References

- [1] Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, 18, 355–387.
- [2] Bock, J. K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology*, 129(2), 177–192.
- [3] Ferreira, V. S. (2003). The persistence of optional complementizer production: Why saying "that" is not saying "that" at all. *Journal of Memory and Language*, 48, 379–398.
- [4] Gries, S. T. (2005). Syntactic priming: A corpus-based approach. *Journal of Psycholinguistic Research*, 34(4), 365–399.
- [5] Jaeger, T. F., Levy, R., Wasow, T., & Orr, D. (2005). Optional *that* is predictable if a relative clause is predictable. AMLaP-2005, Ghent, Belgium.
- [6] Roland, D., Elman, J. L., & Ferreira, V. S. (2005). Why is *that*? Structural prediction and ambiguity resolution in a very large corpus of English sentences. *Cognition*, 98, 245–272.
- [7] Szmrecsányi, B. M. (2005). Language users as creators of habit: A corpus-based analysis of persistence in spoken English. *Corpus Linguistics and Linguistic Theory*, 1, 113–149.

Is frequency a property of phonological forms? Evidence from spontaneous speech

Susanne Gahl (University of Illinois, Urbana-Champaign)

gahl@uiuc.edu

Numerous experimental studies have investigated homophones, with two broad aims: (1) Elucidating the architecture of the language production system: Is there feedback from phonological forms to lemmas, as in interactive models? Do homophones thus activate each others' lemma nodes? (See, e.g., Dell, 1990; cf. also Griffin, 2002.); (2) Understanding the role of experience, as reflected in usage frequency, in language production. For example, Jescheniak & Levelt, 1994 (see also Jurafsky, 2003) argue that lexical frequency is a property of phonological forms, not lemmas.

One piece of relevant evidence has come from pronunciation: High frequency words tend to shorten (cf. Jurafsky et al., 2001). If high-frequency words (e.g., *straight*) do not differ in duration from their low-frequency twins (e.g., *strait*), then this is consistent with models in which frequency is a property of phonological forms, causing low-frequency words to inherit frequency effects from high-frequency twins. On the other hand, if homophones differ in duration, then this supports the view that frequency may be a property of lemmas.

Experimental evidence on homophone duration has been inconclusive (e.g., Geffen & Luszcz, 1997; Wright, 1997; but Guion, 1995, Nygaard et al., 2002, Whalen, 1991). Moreover, this evidence is problematic, since pronunciation may be affected by participants' awareness of homophones in the experiment, and by other artifacts of elicited production.

We examined this issue using the CELEX lexical database (Baayen et al., 1993) and a time-aligned transcript of the Switchboard corpus of 240 hours of American English telephone conversations (Godfrey et al., 1992; Deshmukh et al., 1998). We first extracted all word forms from CELEX that were homophonous with at least one other word. We then extracted the durations of all tokens of these words from Switchboard. Excluding items with non-homophonous homographs, and pairs whose members had identical frequencies in Switchboard, we arrived at a set of 252 usable homophone pairs, represented by a total of 265546 tokens. Comparing the average duration of the lower-frequency member of each pair to that of its higher-frequency counterpart, we found that the lower-frequency words were longer than the high-frequency ones (384.3 vs. 357.3 ms on average; $z = -13.763$, $p < .001$). To check whether this result was unduly influenced by items of very low frequency, which might have elicited unnaturally slow production, we repeated the test while excluding words with fewer than 5 attested tokens. Doing so yielded a similar result as the overall analysis (354.8 vs. 326.8 ms; $z = -3.740$, $p < .001$). A multiple regression analysis revealed frequency, alongside syntactic category, to be a significant predictor of homophone duration.

Since lemma information (e.g., *straight* vs. *strait*) affected durations, our results are consistent with the notion that frequency can be a property of lemmas, not just of phonological forms. We discuss our results within the framework of exemplar-based models of pronunciation (e.g., Johnson, 1997; Pierrehumbert, 2003), which can capture subphonemic detail. At the methodological level, our study illustrates the usefulness of large corpora of spontaneous speech for psycholinguistic research.

Testing the strength of the spurious licensing effect for negative polarity items

Ming Xiang, Brian W. Dillon & Colin Phillips (University of Maryland, College Park)

mxiang@umd.edu

Negative polarity items (NPIs) such as *any* or *ever* are words that are licensed in specific linguistic environments. Generally speaking, NPI-licensors must possess certain semantic / pragmatic properties, but also must be structurally accessible to the NPI by occupying a c-commanding position. In an acceptability judgment task with speeded presentation, Vasishth et al. (2005) showed that German speakers accurately distinguished sentences with a c-commanding negative licensor (1a) from sentences lacking a licensor (1b), but in around 15% of trials accepted sentences with a structurally inaccessible negative element (1c). This 'spurious licensing' effect may reflect a general insensitivity to the structural conditions on NPI licensing, or alternatively may reflect the high surface co-occurrence frequency of the negative determiner with NPIs. We report results from two experiments that show that the spurious licensing effect extends to English, and is not restricted to the most frequent licensors of NPIs.

We tested similar structures to Vasishth et al. (2005), but with three different NPI licensors for the NPI *ever*: *no*, *few* and *only* (2). The NPI *ever* frequently co-occurs with transparently related negative words (e.g., *no*, *not*, and *nobody*). These precede around 30% of instances of *ever*, 12.5% for *no* alone. In contrast, *ever* is preceded by *few* and *only* in only 2.4% and 7.2% of occurrences, respectively (Gigaword Corpus). If the spurious licensing effect reflects the high co-occurrence of negative NPs with *ever*, then the effect should be weaker with other licensors. The three licensors appeared in grammatically accessible positions (2a–c) and grammatically inaccessible positions (2d–f), and in a final condition no potential licensor was present (2g). Twenty-eight sets of 7 items were distributed across 7 lists in a Latin Square design and combined with 84 filler items. An off-line rating task (Experiment 1, N=10) confirmed the importance of a grammatically accessible licensor, and showed only a weak spurious licensing effect (Table). In contrast, an on-line acceptability judgment task using speeded RSVP presentation (400 ms per word) found strong spurious licensing effects (Experiment 2, N=21). Conditions with grammatically accessible licensors were overwhelmingly accepted and sentences containing no licensor were predominantly rejected. However, sentences with grammatically inaccessible licensors were judged acceptable on 42% of trials. These ratings differed significantly from the grammatically accessible conditions ($p < .01$) and from the no licensor condition ($p < .001$). Crucially all three licensors showed the same pattern.

Our results confirm and extend the findings from Vasishth and colleagues, and indicate that they are not restricted to the most frequent licensors for NPIs. There could be two possible interpretations of these results. The spurious licensing effect may reflect a search for a licensor that traverses a complete hierarchical structure but is partially insensitive to structural constraints, as in the ACT-R based account proposed by Vasishth et al.; alternatively it may reflect a search processes that yields false positives because the hierarchical structure is not fully completed during the speeded presentation paradigm. To distinguish these possibilities, we are currently conducting an experiment that manipulates the distance between the licensor and the NPI.

Examples

(1) (These are translated from original German material)

- a. No man who had a beard was *ever* happy.
- b. *A man who had a beard was *ever* happy.
- c. *A man who had no beard was *ever* happy.

(2)

	Expt 1 Offline Rating (1–5 scale)	Average	Expt 2 Speeded Presentation Accuracy Rate	Average
a. No bills that the Democratic senators have supported will <i>ever</i> become law.	4.3		79%	
b. Very few bills that the Democratic senators have supported will <i>ever</i> become law.	4.0	4.1	81%	84%
c. Only three bills that the Democratic senators have supported will <i>ever</i> become law.	4.1		93%	
d. * The bills that no Democratic senators have supported will <i>ever</i> become law.	2.1		62%	
e. * The bills that very few Democratic senators have supported will <i>ever</i> become law.	2.2		60%	
f. * The bills that only three Democratic senators have supported will <i>ever</i> become law.	2.2	2.2	52%	58%
g. * The bills that the democratic senators had supported will <i>ever</i> become law.	1.7	1.7	81%	81%

Reference

Vasishth, S., Drenhaus, H., Saddy, D., & Lewis, R. L. (2005). Processing negative polarity. Poster presented at the 18th Annual CUNY Conference on Human Sentence Processing, Tucson AZ.

Parsing quantifiers in object position

Jason Varvoutis & Martin Hackl (Pomona College)

Martin.Hackl@pomona.edu

1. Introduction

The problem of quantifiers in object position is a well-known case of syntax semantics mismatch (Montague 1973) that can be summarized as follows: Even though quantifiers cannot be analyzed as referring expressions and internal argument positions of predicates are reserved for such expressions, quantifiers seem to occur freely in those positions (e.g., *John likes every girl.*). Various solutions to this problem exist (e.g., type-shifting, quantifier raising, continuations). All maintain that semantic composition is more complex for quantifiers in object positions than for quantifiers in subject position. Whether this increased complexity affects real time processing of sentence structure has not been investigated as far as we know. Positive evidence to this effect would therefore be an important contribution to the parsing literature. More specifically, it would show that purely formal semantic complexity affects sentence processing even though — and unlike the cases of semantic coercion investigated in, e.g., McEree et al. (2001, 2002, 2005) — there is no concurrent change in meaning.

2. NP/S Attachment Ambiguity Resolution with Quantifiers

To study whether the parser "knows" about the semantic complexity incurred by quantifiers in object position we ask whether semantic properties of DPs (quantificational/referring) are a factor in the resolution of temporary attachment ambiguities. We employ sentences that are locally ambiguous between construing a DP as the object of the preceding verb or the subject of an embedded clause (NP/S ambiguity; cf. Trueswell, et al. 1993, etc.). That is, as exemplified in (1), we created a situation in which the parser cannot know if the DP/QNP is the object of 'remembered' or the subject of the sentential complement until disambiguation occurs on 'who' or 'was'. If the semantic complexity incurred by a quantifier in object position affects the parser object disambiguation ('who') should be dispreferred over subject disambiguation ('was') in the case of *every NP*. Assuming that definite descriptions can be interpreted as referring expressions (in the default case), no such effect is expected. That is, if the parser "knows" about quantifiers in object position we expect an interaction between determiner type (quantifier / definite determiner) and attachment type (NP/S) in the post disambiguation area.

To control for possible interference of the matrix verb which could create a spurious interaction we chose only S-biased matrix verbs (Trueswell et al., 2003, etc.). Averaging residual reading times over 20 subjects, we obtained two effects:

1. Reading times on the noun immediately following 'the' or 'every' show a main effect of determiner ($p=.036$) such that 'every NP' takes longer than 'the NP'. This shows that the semantic difference between 'the' and 'every' is reflected in real time processing.
2. There is a significant interaction on the first ($p=.045$) word after disambiguation as well as in the region of the word of disambiguation to the third word after that ($p=.01$) between determiner type and attachment type indicating that quantifiers in object position are more difficult for the parser than quantifiers in subject position or definite descriptions in object or subject position.

Example

- (1) The nun remembered the / every child (who) was abused and malnourished.

On-line interpretation of scalar quantifiers: Insight into the semantic-pragmatic interface

Yi Ting Huang & Jesse Snedeker (Harvard University)

huang@wjh.harvard.edu

The relationship between semantic and pragmatic interpretations has been a perennial puzzle in psycholinguistics. While many aspects of utterances are tightly linked to word meaning and syntactic structure, other facets are clearly added by context-sensitive, inferential interpretative processes. General agreement about the existence of these levels of interpretation contrasts with controversy in the field over exact boundaries between these representations and their relationship in real-time processing. One possibility is that semantics strictly precedes pragmatic procedures which are only computed over fully-formed semantic representations. Alternately, pragmatic interpretations could occur simultaneously with semantics to flesh out an underspecified representation.

Past research has generally pursued these issues by exploring how contextual information influences reaction times for sentence interpretations (Frisson & Pickering, 1999; Sedivy et al., 1999). These studies demonstrate that pragmatic processing is rapid, often beginning before phrasal completion. However, they leave open the question of whether semantic interpretations of words precede pragmatic ones. The current experiments address this question by examining a relatively well understood test case from linguistics. Horn (1972) noted that scalar quantifiers like *some* have distinct readings that reflect semantic and pragmatic levels of interpretation. Semantically, it does not possess an upper-bound and is compatible with the total quantity, e.g., "*Did you win some (AND POSSIBLY ALL) of the games?*" This differs from typical interpretations involving a pragmatic implicature which adds an upper-bound to exclude the total set, e.g., "*I won some (BUT NOT ALL) of the games.*" The dual readings of *some* create an ideal situation for exploring the relationship between semantic and pragmatic processing because the meaning assigned at each level can correspond to different arrays.

We investigated how these interpretations unfold over the course of on-line speech comprehension using the visual-world eye-tracking paradigm. Participants were asked to "*Point to the girl that has some of the socks*" when presented with a girl with 2 out of 4 socks (NOT ALL) and another girl with 3 of 3 soccer balls (ALL). These critical trials were semantically ambiguous since the meaning of *some* is consistent with both interpretations prior to the final phonological disambiguation of the noun. These were compared to control trials where participants heard quantifiers with semantically specified upper-bounds (*two, three, all*).

Experiment 1 revealed that participants preferred to look at the target picture after hearing *two, three, and all*, indicating early disambiguation based on lexical semantics of the quantifier. In contrast, after hearing *some*, participants initially looked at both girls with no preference for the target in the ambiguous region. This resulted in a significant difference between *some* and the other quantifiers ($p < .01$). In Experiment 2, we found a reliable target preference for the target 400 ms after quantifier onset for *two, three* and *all*, but no reliable preference for *some* until 800ms after quantifier onset (see Table 1). Nevertheless, target-preference for *some* trials preceded the phonological disambiguation of the noun, suggesting that participants generated the implicature. These results demonstrate that lexical semantics guide interpretation before even the most robust pragmatic inferences are calculated.

Table 1. In Experiment 2, proportions of looks to target were calculated over time from onset of the quantifier. Analysis of a main effect of quantifier type is listed over each interval.

	Time from onset of the quantifier				
	200 ms	400 ms	600 ms	800 ms	1000 ms
Some	0.39	0.41	0.48	0.61	0.84
All	0.57	0.63	0.78	0.82	0.93
Two	0.52	0.67	0.79	0.79	0.88
Three	0.56	0.70	0.79	0.86	0.95
p-value	.08	.01	.01	.01	.01

Effects of a word's status as a predictable phrasal head on lexical decision and eye movements

Adrian Staub & Charles Clifton, Jr. (University of Massachusetts Amherst)

astaub@psych.umass.edu

Wright and Garrett (1984) presented data suggesting that readers are speeded in processing a word whose syntactic category is predictable. Their subjects made lexical decision responses to target words that were presented following context sentences. Lexical decision RT was faster when the target was a predictable phrasal head than when this word was a legal, but unpredictable, continuation of the sentence. We performed a series of five experiments that extended Wright and Garrett's basic finding, while at the same time ruling out several artifactual explanations. In Experiments 1, 2, and 4, participants read context fragments that were presented word-by-word in an RSVP paradigm, followed by a lexical decision target in uppercase letters. In filler trials the sentences continued to the end, after which participants were tested for comprehension. In Experiments 3 and 5, readers' eye movements were monitored as they read full sentences.

Experiments 1–3 used sentences like (1a–b) below, in which either an adjective-noun sequence (*fancy furniture*) or a noun-noun compound (*porch furniture*) appeared in subject position in an embedded clause. In Experiment 1, the lexical decision target was the word that varied between conditions (*fancy / porch*). RT was significantly faster to nouns than to adjectives, though these words did not differ in length, frequency, or cloze probability (which was near zero in both conditions). In Experiment 2, the target was the subsequent noun (*furniture*), and RT was significantly faster when this word followed an adjective than when it followed a noun. Again, cloze probability was near zero in both conditions. In Experiment 3, readers' eye movements were monitored as they read full sentences like (1a–b). Gaze duration was significantly shorter on the nominal modifier than on the adjectival modifier, replicating the result of Experiment 1. Reading time on the head noun did not differ significantly between conditions, which we attribute to spillover effects.

Experiments 4 and 5 used sentences like (2a–b) below to explore whether adjective processing is speeded following a degree adverb, which renders an adjective a predictable phrasal head. In Experiment 4, the critical adjective (*flexible*) was the lexical decision target. RT was faster in the degree adverb condition, and there was no difference between control conditions in which the target was a nonword. In Experiment 5, readers' eye movements were monitored as they read full sentences like (2a–b). In the degree adverb condition, gaze duration was numerically shorter on the adjective, and significantly shorter on the subsequent noun. Again, this pattern is interpreted as reflecting spillover processing.

The results confirm that lexical processing is affected by a word's status as a predictable phrasal head. This is likely to be a relatively late effect, i.e., not an effect on lexical access itself, but an effect on syntactic integration. We interpret the results with respect to top-down parsing theories, according to which the parser posits predictable structure before encountering the corresponding lexical input.

Examples

- (1) a. The supervisor decided that the fancy furniture would no longer be produced.
b. The supervisor decided that the porch furniture would no longer be produced.
- (2) a. The mechanic used some very flexible plastic to fix the problem.
b. The auto mechanic used some flexible plastic to fix the problem.

References

Wright, B., & Garrett, M. (1984). Lexical decision in sentences: Effects of syntactic structure. *Memory & Cognition*, 12, 31–45.

Relation between temporal adverbs and verb morphology in agrammatic aphasia

Yasmeen Faroqi-Shah (University of Maryland, College Park)

yshah@hesp.umd.edu

Aphasia is an impairment in producing and/or comprehending language after acute brain injury. Agrammatic speech is a frequent occurrence in aphasia and is characterized by ill-formed sentences and errors of grammatical morphology, particularly verb inflections. Typical verb inflection errors include substitutions such as, *Yesterday she speaks/ speaking at the meeting*. The cause of verb morphology errors in aphasia is unknown. Since verb inflection errors are cross-linguistically most frequent with tense morphology (Friedman & Grodzinsky, 1997) and given the typical pattern of substitutions (see example above), Faroqi-Shah and Thompson (2003) argued that verb inflection errors arise from a morphosemantic difficulty in encoding temporal information and/or retrieving verb forms that match the intended tense.

In this study, two questions that arise from Faroqi-Shah and Thompson's proposal are addressed: 1) Are the two morpho-semantic processes, i.e., encoding temporal information and retrieving corresponding verb form, separable? And if so, which of these is the locus of deficit in morphological agrammatism? A deficit in encoding temporal information will affect both morphological and lexical elements that convey temporal information. That is, there will be parallel deficits in production of temporal adverbs and verb morphology. In contrast, a verb form retrieval deficit will not impair the production of temporal adverbs. 2) Do adverb processing difficulties (e.g., *yesterday = past tense*) contribute to inflectional substitutions such as in the above example? There is little research on the production of adverbs in agrammatic aphasia. A general deficit with adverbs will affect the production of temporal adverbs as well as other adverbs that do not convey temporal information.

Procedures. A multiple choice sentence completion task was used with three conditions: verb inflections, temporal adverbs, and other adverbs. Stimuli used to elicit verb inflections consisted of temporal cues within adverbial clauses:

- (1) *After Mary moved the sofa, she _____ her back.* Response options: *sprained, sprains, will sprain*

In stimuli used to elicit temporal adverbs, temporal information was conveyed via verb morphology:

- (2) _____ *the janitor sweeps the dirty floor.* Response options: *Everyday, Yesterday, Last year*

Other adverb stimuli elicited a variety of manner, frequency, degree and comment adverbs:

- (3) *Mary _____ called 911 after the accident.* Response options: *quietly, slowly, frantically*

So far, five patients with agrammatic aphasia have been tested (total N = 10).

Results & Discussion. Mean accuracy scores were 38.1% (SD = 12.8), 80.4% (SD = 13.4), and 90% (SD = 5.8) for verb morphology, temporal adverbs and other adverbs respectively. For all the agrammatic aphasic patients, verb morphology scores were significantly lower than temporal adverb and other adverb scores (Fisher Exact $p < .05$) and there was no difference between the accuracies of temporal and other adverbs (Fisher Exact $p = 1$). These results suggest a unidirectional deficit in agrammatic individuals: they could successfully select temporal adverbs from cues in verb morphology, but had difficulty in selecting verb morphology from temporal cues. That is, the morphosemantic deficit that results in verb inflection errors is restricted to retrieval of verb form rather than generalized temporal encoding.

References

- Faroqi-Shah, Y., & Thompson, C. K. (2003). Verbal inflections in agrammatic aphasia: A multilevel analysis. Platform presentation: *American Speech-Language and Hearing Association*, Chicago IL.
- Friedman, N., & Grodzinsky, Y. (1997). Tense and agreement in agrammatic production: Pruning the syntactic tree. *Brain and Language*, 56, 397–425.

Speakers' control over leaking private information

Liane Wardlow Lane, Michelle Groisman & Victor S. Ferreira (University of California, San Diego)

lwardlow@ucsd.edu

Past research demonstrates that speakers sometimes make references to privileged objects (objects known only to them) when naming mutually visible objects (Horton & Keysar, 1996; Nadig & Sedivy, 2002; Wardlow & Ferreira, 2003). For example, Wardlow and Ferreira (2003) report a task where speakers and addressees were presented with four cards each depicting a simple object. Both could see the same three objects (i.e., a circle, a square, and a triangle), but the speaker could see an additional, privileged object (a smaller triangle). Speakers were asked to identify one of the mutually visible objects (the *target*) for the addressee. When asked to identify the triangle, speakers *should* have said "triangle." However, they often said "large triangle", as if they failed to account for perspective differences. Interestingly, such utterances serve to implicitly leak extra information. Here, "large triangle" conveys that the speaker can also see another, smaller triangle. But can speakers avoid communicating implicit information when doing so conflicts with their goals?

We used a referential communication task like that described above. On *test* trials, the privileged object was the same as the target object but differed in size, whereas on *control* trials, the privileged object was distinct. In the *baseline* block, speakers were simply asked to name a target. In *conceal* blocks, participants were given additional instructions that encouraged speakers to hide the identity of the foil when identifying the target. Specifically, after addressees selected the target, they could guess the identity of the privileged object. Speakers and addressees kept scores; a correct guess gave addressees an additional point. Thus, speakers were provided with both incentive and instruction to conceal the identity of the privileged object. If speakers can control leaking information, then the conceal instruction should reduce modifier use relative to baseline performance.

Results showed that on test trials, speakers used modifying adjectives *more* in the conceal condition (14.4%) than in the baseline condition (5.4%). Speakers rarely used modifying adjectives in the control conditions (1.4% and 0.5%). Thus, the instruction to conceal privileged information made speakers refer to it even more; this is likely because the instruction to conceal privileged objects served to make them highly salient, and the production system had a difficult time blocking the intrusion of such information. These results localize perspective-taking errors to a stage of processing, grammatical encoding, that is outside speakers' executive control. Additionally, the results suggest not only that leaked information may be information speakers want to keep private, but that attempts to conceal it might make its leakage even more likely. If so, these results are likely to be relevant to interactions involving everything from interpersonal interactions to adversarial negotiation.

References

- Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition*, 59(1), 91–117.
- Nadig, A. S., & Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science*, 13(4), 329–336.
- Wardlow, L., & Ferreira, V. S. (2003). Finding common ground: How do speakers block privileged information? The AMLaP Conference, Glasgow, Scotland.