

The role of veridicality and factivity in clause selection

Aaron Steven White
University of Rochester

Kyle Rawlins
Johns Hopkins University

Introduction. Which of a verb’s semantic properties are active in determining its selectional properties? This question has received surprisingly different answers for nominal and clausal selection. Whereas accounts of nominal selection tend to be stated in terms of event structure (see Levin & Rappaport Hovav 2005 for a review), accounts of clausal selection tend to be stated in terms of a distinct set of semantic-pragmatic properties such as ‘representationality’ (see White 2015 for a review).

Contribution. We investigate two properties that have been claimed to be important in determining clausal selection: *veridicality* and *factivity*. Using a new dataset of veridicality judgments for all English verbs that embed declarative clauses, we show that these properties do not predict any selectional properties above and beyond what can be predicted from verb class alone.

Background. Hintikka (1975) claims that a verb embeds both declaratives and interrogatives—i.e. is *responsive* (Lahiri 2002)—iff it is *factive*—i.e. presupposes the content of its embedded clause. This generalization covers a wide range of responsive predicates—including cognitive factives (e.g. *know*, *realize*, *discover*, *find out*) and emotive factives (e.g. *love*, *hate*, *bother*)—though it has various well-known exceptions. For instance, Egré (2008) shows that many non-factive *veridical* predicates—i.e. predicates that merely entail the content of their embedded clause (e.g. *prove*)—are also responsive. Assuming that all factives are veridical (and thus that factives’ presuppositions are entailed), Egré suggests a revised generalization: a verb has the responsive selectional pattern iff it is veridical.

Even this revised generalization has a large class of exceptions that Egré himself acknowledges: many communicative predicates (e.g. *say*, *tell*) are responsive yet (apparently) nonveridical. This class of counterexamples can be addressed in various ways, but importantly this approach relies on a notion of verb class to constrain the generalization—e.g. responsive communicatives might be ambiguous between veridical and nonveridical variants (Spector & Egré 2015).

A question that arises here is whether reference to veridicality is necessary above and beyond knowledge of a particular verb class. Recent work suggests that, even in the domain of noncommunicatives, there are systematic exceptions to the apparent correlation between veridicality and responsivity (White & Rawlins 2017) that appear to be dependent on event structure. For instance, there is a small subclass of noncommunicative predicates that are nonveridical but responsive: change-of-state predicates, such as *decide* and *choose*. Nonetheless, it could be that there are imperfect-yet-nontrivial correlations between veridicality/factivity and responsivity that are not reducible to event structure.

We approach this problem at a large scale by building on the MegaAttitude data set of White & Rawlins (2016), which contains acceptability judgments for effectively every English clause-embedding verb in 50 different subcategorization frames. These frames instantiate configurations of syntactic features that have been suggested to be correlated with some aspect of clause-embedding verb semantics: COMPLEMENTIZER, EMBEDDED TENSE, MATRIX OBJECT, MATRIX PP, EMBEDDED SUBJECT, and PASSIVIZATION. Further, to avoid typicality effects, these frames were constructed to contain as little lexical content as possible besides the verb at hand. This was done by ensuring that all DP arguments were indefinite pronouns and all verbs besides the one being tested were either *do* or *happen*. Example (1) shows two examples of the kind of sentences participants judged.

- (1) a. Someone thought something happened. b. Someone was told to do something.

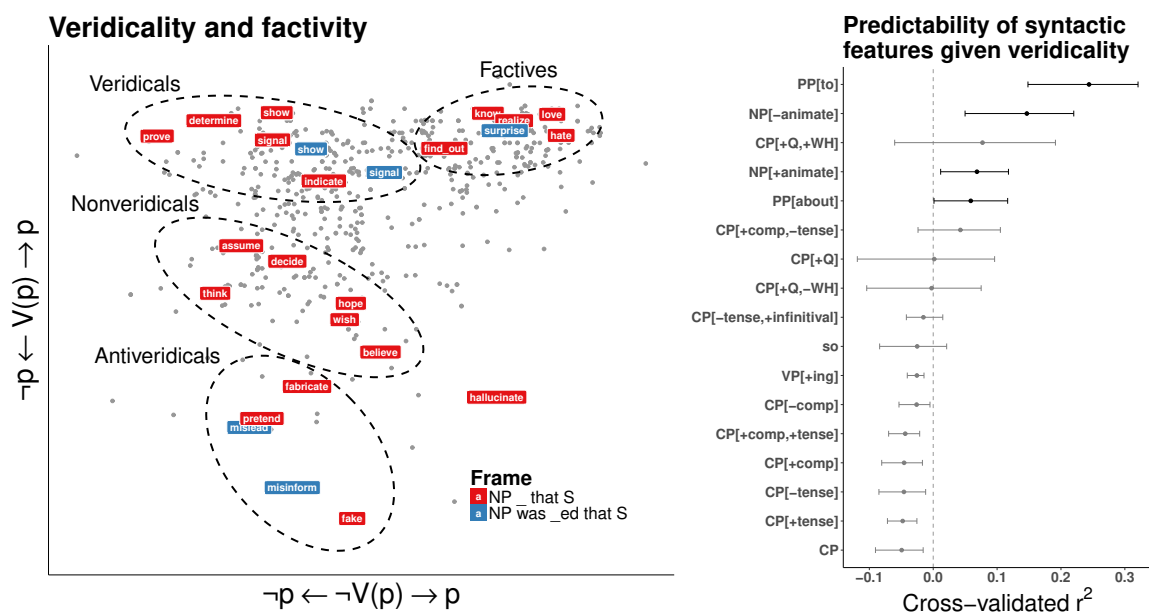
We use the MegaAttitude dataset to obtain a measure of the selectional relationship between a clause-embedding verb and each of the features listed above, and we use the general method of constructing low content items, such as those in (1), to avoid typicality effects for our veridicality experiment as well.

Experiment. We selected 517 verbs from MegaAttitude based on their acceptability in the [NP _ that S] and [NP was _ed that S] frames—on average, a rating of 4 out of 7 or better. This yielded a sample of 348 verbs that only occurred in the active frame, 142 that only occurred in the passive frame, and 27 that occurred in both. These verbs were then slotted into contexts such as those in (2) based on which set they fell into. This yielded 1,088 items, which were randomly partitioned into 16 lists of 68.

- (2) a. Someone {thought, didn’t think} that a particular thing happened.
b. Someone {was, wasn’t} told that a particular thing happened.

Given sentences such as those in (2), the task was to answer the question *did that thing happen?* There were three response options: *yes*, *maybe* or *maybe not*, and *no* (cf. Karttunen et al. 2014).

We recruited 160 unique participants through Amazon’s Mechanical Turk, yielding 10 ratings per item given by 10 different participants. Item order for each list was randomized for each participant. To



control for differences in participant’s use of the response scale, we applied a rdit-scoring transformation to their responses prior to analysis (see Agresti 2013).

Results. The left figure plots the normalized responses for contexts with negative matrix polarity (x -axis) against those for contexts with positive matrix polarity (y -axis). A verb being toward the top means that it entails its embedded clause, while a verb being towards the bottom means that it entails the negation of its embedded clause. A verb being toward the right means that its embedded clause content projects through negation, while a verb being towards the left means that the negation of its embedded clause content projects through negation. We see that our experiment successfully distinguishes between factives (upper right), veridicals (upper left), nonveridicals (middle), and antiveridicals (lower left).

Analysis. To ask whether veridicality or factivity are predictive of any syntactic distinction without reference to verb class, we first attempt to predict the acceptability judgments from MegaAttitude given our data. We obtain a measure of verbs’ selection for particular syntactic features by taking the maximum acceptability of that verb for any frame that instantiates that feature. This means that, if a verb is good with any frame that instantiates that feature, it will be marked as good with that feature.

For each syntactic feature, we predict the selection measure using a LASSO regression with predictors for the normalized veridicality judgments with positive and negative matrix polarity as well as their interaction. This regression is submitted to a 10-fold nested cross-validation. The mean cross-validated (CV) r^2 is plotted in the right figure with 95% CIs generated by a non-parametric bootstrap. A particular feature is reliably predictable if the error bar for a positive CV r^2 does not cross zero.

We find significant correlation for PP[to], NP[-animate], NP[+animate], and PP[about]. Interestingly, no clausal complement features are reliably predictable, though we see that the CP[+Q,+WH] features has the third highest predictable overall. The fact that the predictability of CP[+Q,+WH] shows such high variability in the cross-validation may indicate why previous authors mistakenly took veridicality/factivity to be correlated with responsibility: if one looks at a small enough subset of verbs, spurious correlations can arise, but these correlation’s generalizability will be poor.

We hypothesize that the correlations that we do find are driven exclusively by verb class—specifically, whether the verb is an emotive or communicative. To test this, we residualize our veridicality/factivity predictors by whether each verb is emotive or communicative—thus effectively removing whatever information about those classes is encoded in the veridicality judgments. We then we rerun the cross-validation using these residualized predictors. We find that after residualization, no syntactic features are reliably predictable given veridicality/factivity. This suggests that verb class—not veridicality or factivity—is the real predictor of these syntactic features.

Conclusion. Our findings for veridicality and factivity are just one instance of what we believe will turn out to be a general pattern: semanticopragmatic properties that can only hold of clause-embedding verbs, such as veridicality and factivity, are not the sort of lexical semantic property that selection traffics in. Rather, we should seek a theory of selection that makes reference to properties that are instantiated across the lexicon more generally, such as stativity, dynamicity, and telicity. In future work, we plan to build datasets similar to the one presented here for both event structural properties and other non-event structural properties that have been proposed to be relevant to clausal selection, such as *neg-raising*. We hypothesize that, using explicit statistical model comparison, we will find that models that predict syntax based on event structural properties will outperform ones based on non-event structural properties.