

# Research Statement

Aaron Steven White ([aaronstevenwhite.io](http://aaronstevenwhite.io))

## Overarching Goals

The human ability to convey information about possible past, present, and future configurations of things in the world is undergirded by systematic relationships between linguistic expressions and ontological/conceptual categories of arbitrary granularity. My research investigates (a) which ontological/conceptual categories can be related to which sorts of expressions in a language, (b) what those relationships look like, and (c) how they can be leveraged for building better natural language understanding (NLU) systems.

## Research Program

My research program has two components: the design, implementation, and deployment of (i) precision instruments for measuring the distributional and inferential properties of linguistic expressions; and (ii) scalable systems for synthesizing those data to provide both scientific insights in natural language semantics and modular components of NLU systems. My current approach consists in two complementary projects: the **MegaAttitude Project** and the **Decompositional Semantics Initiative**.

### MegaAttitude ([megaattitude.io](http://megaattitude.io))

The MegaAttitude Project aims to (a) test hypothesized generalizations about the relationship between the inferential characteristics and morphosyntactic distributions of predicates that combine with subordinate clauses and (b) develop computational models for discovering such generalizations and the underlying lexical semantic properties that drive them by exhaustively cataloguing the distributional and inferential characteristics of all such predicates in English.

**Key papers:** White and Rawlins 2016, 2018a,b, 2020; White 2019, 2021; Moon and White 2020; An and White 2020; Kim and White 2021; Kane et al. 2021

### Decomp ([decomp.io](http://decomp.io))

The Decompositional Semantics Initiative (Decomp) aims to capture a rich array of theoretically motivated semantic properties in context by annotating genre-diverse corpora. These annotations have been used both for testing hypotheses about the nature of thematic roles and event structure classes as well as for developing state-of-the-art syntactic and semantic parsers and natural language inference systems. I recently gave an interview about Decomp on the [NLP Highlights](#) podcast.

**Key papers:** White et al. 2016, 2017b, 2018b, 2020; Vashishtha et al. 2019; Govindarajan et al. 2019; Stengel-Eskin et al. 2020, 2021; Gantt et al. 2021

## Future Directions

It is becoming increasingly clear that integrating data from targeted behavioral experiments with evidence from attested language use is integral for the development of broad-coverage syntactic and semantic theories. My goal in the coming years is to build on models of behavioral data developed under the MegaAttitude Project and annotated corpus data developed under Decomp to construct unified computational models that synthesize the two forms of data within a theoretically informed framework. I recently laid out a blueprint for how this might be done at the sentence level (Kim and White 2021), which I have begun to expand to the discourse/document level (Gantt et al. 2021).

## Background

**Predicates** that combine with subordinate **clauses**—e.g. cognitive (*think, remember*), desiderative (*want, wish*), emotive (*love, surprise*), communicative (*say, tell*), and aspectual (*start, stop*) predicates, among others—are useful for studying the relationship between ontological/conceptual categories and linguistic expressions due to the complexity of their distributional and inferential characteristics.

Jo **forgot** { **that Bo left<sub>1</sub>**, **to leave<sub>2</sub>**, **leaving<sub>3</sub>** }.

For instance, when *forget* is combined with a finite declarative complement (1), we get an inference that the content of the complement is true—that Bo did in fact leave—but when it is combined with an infinitival complement (2), we get an inference that the content is false—that Jo did not in fact leave. The latter inference is reversed if we replace *forget* with *remember*, yet other inferences remain the same: when combined with infinitival complements (2), both *remember* and *forget*, we get an obligation inference—that Jo was supposed or planned to leave—not present with other nonfinite complements (3).

## Overarching Questions

These inferential patterns raise questions about (i) which components of an expression—the predicate, subordinate clause, or both in interaction—carry the information on which the inference is based; (ii) whether that information is idiosyncratic to a predicate or general to a class and, insofar as they are general, how complex the inference patterns associated with any such class are; and (iii) whether and how such classes correlate with syntactic distribution. These questions can only be answered definitively by collecting acceptability and inference judgments for a broad sample of the lexicon.

## Approach and Main Findings

I have collected large-scale datasets capturing the acceptability of over 1,000 predicates in 50 different syntactic contexts (White and Rawlins 2016, 2020; An and White 2020) as well as a wide variety of inference types (White and Rawlins 2018b; White et al. 2018b Moon and White 2020; An and White 2020; Kane et al. 2021). I have found that, while the source of the information relevant to triggering an inference varies across inference type—with certain types of inferences being a product of both a predicate and its subordinate clause (An and White 2020) and others being the product of only one (Moon and White 2020)—the inferences themselves are highly correlated, coalescing into complex patterns associated with conceptually coherent clusters that predict syntactic distribution to a surprising degree (Kane et al. 2021). These correlations are especially surprising in light of the fact that the individual inferences from which these patterns are composed do not themselves correlate strongly with distribution (White 2021).

## Upshots and Open Questions

These findings are consistent with lexical semantic theories built for other areas of the lexicon (White and Rawlins 2018a) as well as theories of word learning that put much of the burden on the use of syntactic cues, where tight correlations among inferences may be crucial (White 2015, White et al. 2018c). A major remaining question is the extent to which these classes are further decomposable into a small set of representational components that explain both the inference patterns associated with a class and the syntactic distribution of predicates in that class. I am currently pursuing these questions using syntactic and semantic grammar induction systems (Kim and White 2021).

## Background

Some inferences conventionally associated with particular linguistic expressions (lexical items, syntactic structures, or their combination) are subject to contextual modulation. For instance, if someone were to say out of the blue *Jo recently stopped smoking*, we might infer that Jo used to smoke. This inference is plausibly triggered by the fact that one cannot stop something they were not doing in the first place. But if this expression were instead embedded in a conversational context like *I don't know if Jo used to smoke, but if she recently stopped smoking, she's eligible for lower cost insurance*, we do not make a similar inference. This sort of contextual modulation suggests the usefulness of annotating for inferences in the context of use. Such annotations can support both theory development and the development of natural language understanding (NLU) systems.

## Challenges for Existing Frameworks

Traditional frameworks for capturing inferences in context generally define complex, often exclusive category systems that require highly trained annotators to build. And in spite of their high quality for the cases they are designed to handle, these frameworks can be brittle to cases that (i) deviate from prototypical instances of a category; (ii) are equally good instances of multiple categories; or (iii) fall under a category that was erroneously excluded from the framework's annotation ontology.

## Approach

I have proposed an alternative approach to semantic annotation rooted in a long tradition of theoretical approaches to lexical semantics that addresses these issues by approaching semantic annotation via many simple independent questions about words or phrases in context (White et al. 2016). I have deployed this framework to capture a wide variety of inference types (Rudinger et al. 2018; Vashishtha et al. 2019; Govindarajan et al. 2019; Gantt et al. 2021), which I have combined into a unified syntactic and semantic graphbank with an easy to use toolkit (White et al. 2020).

## Developing Semantic Theories

One use I have put this dataset to is evaluating and generalizing theories of the relationship between conceptual/ontological categories and distributional properties. In this vein, I have investigated which inferences about the relationship between an event and its participants are relevant in determining the grammatical position (e.g. subject, direct object, etc.) of a linguistic expression of those participants relative to a predicate describing that event (White et al. 2017b). I have recently generalized this approach to investigate the categories into which event descriptions and relations among events fall (Gantt et al. 2021). This latter work opens up new avenues for investigating how languages express information about complex events.

## Developing NLU Systems

Another use I have put this dataset to is evaluating and developing natural language understanding systems. In terms of evaluation, I have used this dataset to probe which inferences generic natural language inference systems capture (White et al. 2017c, 2018b; Poliak et al. 2018; Vashishtha et al. 2020). In terms of development, I have used it to build systems for parsing text into syntactic and semantic graphs (Stengel-Eskin et al. 2020). These parsers obtain state-of-the-art or near state-of-the-art performance on standard multilingual syntactic parsing tasks, suggesting exciting directions for leveraging semantic information for syntactic parsing across a variety of languages (Stengel-Eskin et al. 2021).