

Inference Patterns with Intensional Adjectives

James Pustejovsky

Brandeis University

jamesp@cs.brandeis.edu

Abstract

In this paper we report on an ongoing multi-institution effort to encode inferential patterns associated with adjective modification in English. We focus here on a subset of intensional adjectives typically referred to as “non-subjective” predicates. This class includes adjectives such as *alleged*, *supposed*, *so-called*, and related modally subordinating predicates. We discuss the initial results of corpus-based investigations to discriminate the patterns of inference associated with these adjectives. Based on these studies, we have created an initial annotation specification that we are using to create a corpus of adjective-related inferences in English.

1 Introduction

One of the primary goals for linguistic annotation projects is the explicit representation of the syntactic and semantic information necessary for the creation of useful and meaningful inferential structures. In this brief note, we report on a multi-institution effort underway to identify and model the inferential patterns associated with three diverse classes of adjectives in English. This research combines the efforts of Princeton University (C. Fellbaum), Stanford University (A. Zaenen and L. Karttunen), and Brandeis University (the present author).

Adjectives can be divided into different classes, depending on what dimensions of analysis are being used. Classic semantic field analysis (cf. (Dixon, 1991; Lyons, 1977; Raskin and Nirenburg, 1995)) categorizes the attributes denoted by adjectives according to a thematic organization, centered around

a human frame-of-reference, as lexically encoded in the language, such the following classes:¹ DIMENSION, PHYSICAL PROPERTY, COLOR, EMOTIONS, TEMPORAL SPATIAL VALUE, MANNER.

As intuitive as these classes might be for organizing aspects of the lexis of a language, they fail to provide a coherent guide to the inferential patterns associated with adjectival modification. An alternative approach is to adopt a conceptually conservative but more formally descriptive and operational distinction, one which groups adjectives into inferential classes. (Amoia and Gardent, 2006) and (Amoia et al., 2008), following (Kamp, 1975) and (Kamp and Partee, 1995), make just such a move, adopting a four class distinction based on inferential properties of the adjective, as illustrated below:

- (1) In the construction, [A N], A can be classed as:
 - a. INTERSECTIVE: the object described is both A and N.
 - b. SUBJECTIVE: the object described is A relative to the set of N, but not independent of N.
 - c. PRIVATIVE: the object described is not an N, by virtue of A.
 - d. NON-SUBJECTIVE: there is epistemic uncertainty whether the object is N.

These constructions constitute patterns that license specific inferences associated with classes of adjectives, and can be exploited in the context of text-based inference systems, such as the RTE (Amoia and Gardent, 2006). This classification, however, is both too broadly defined to model the finer inferential distinctions within each class, and too narrow

¹It should be noted that (Raskin and Nirenburg, 1995), however, also discuss inferential patterns for distinct classes.

to include the behavior of other adjective classes, in particular, those taking clausal complements. For these reasons, we have chosen to study three different classes of adjectives that require refinements and additions to the inference patterns given above. These classes are:

- (2) a. Scalar adjectives: both dimensional (*big*, *small*) and evaluative (*happy*, *pretty*) scalars have been categorized as subsective adjectives;
- b. Adjectives with clausal complements: adjectives such as *annoying* and *nice*, when governing clausal complements, do not fit nicely into any of the above classes;
- c. Intensional adjectives: adjectives such as *alleged* and *supposed* are non-subsective, but in complex ways that are dependent on the semantics of the nominal head.

Concerning the third adjective class, the intensional adjectives, the effect of modifying the nominal head is the introduction of epistemic uncertainty regarding the description.

- (3) *T*: The police arrested the **alleged** criminal.

H: A criminal was arrested.

Hence, this inference would be false. Now consider the pair below:

- (4) *T*: Archeologists discovered an **alleged** paleolithic stone tool.

H: A stone tool was discovered.

This inference is legitimate because the epistemic scope of the adjective *alleged* is the adjective *paleolithic*, and not the nominal head itself. In the next section, we look at the behavior of the non-subsective intensional adjectives in more detail, and see that there is a more nuanced, but still systematic, pattern at work.

2 Intensional Adjective Behavior

Recall that intersective adjectives such as *carnivorous* have the following behavior:

$$(5) \|A N\| = \|A\| \cap \|N\|$$

Subsective adjectives, on the other hand, such as *big*, can be modeled as follows:

$$(6) \|A N\| \subseteq \|N\|$$

The intensional adjectives can be split into privatives and non-subsective. Privatives, such as *fake* or *pretend*, can be analyzed as follows:

$$(7) \|A N\| \cap \|N\| = \emptyset$$

Intensional non-subsective adjectives introduce an epistemic uncertainty for the elements within their scope. Examples of this class include *alleged*, *supposed*, and *presumed*, and they call into question some predicative property of the nouns they modify. Following (Kamp and Partee, 1995), no informative inference is associated with this construction:

- (8) a. $[A N]$ (alleged criminal)

b. $\neq N$

However, contrary to what is claimed in (Amoia and Gardent, 2006), non-subsective adjectives do appear to license specific inferences when examined in a broader context than the $[A N]$ construction usually studied. From preliminary corpus studies of this class², several distinct patterns of inference emerge. While the typical resulting composition entails uncertainty of whether the nominal head belongs to the mentioned sortal, (9a) below, there are many contexts where the epistemic scope is reduced to a modification or additional attribution of the nominal head, as shown in (9b).

- (9) a. The **alleged criminal** fled the country.

b. Archeologists discovered an **alleged paleolithic tool**.

In Example (9a), the adjective *alleged* calls into question the predicative property of ‘criminality’ of the *criminal*. When a predicative property is called into question by adjectives of this class, are there any systematic inferences to be made about the semantic field? E.g., is the semantic field still guaranteed to be some hypernym of *criminal*? Even if the individual does not belong to the set of “criminals”, it does still seem to belong to the set of “persons”. In example (9b), contrastively, at least under one interpretation, it is whether the *tool* is *paleolithic* or not that is called into question: i.e., the object belongs to the set of “tools” regardless if it is truly *paleolithic* or not.³ This inference is schematically represented below.

²The initial corpus has been collected from directed CQL queries over two Sketch Engine corpora, Ententen12 and BNC. Three sentence “snippets” have been compiled from this source.

³One reviewer has correctly pointed out that this inference still appears too weak to capture the intended interpretation.

- (10) Given the construction $[A_{int} N]$, where A_{int} is *alleged*, ..., then:
- $[A_{int} N] \not\equiv N$
 - $[A_{int} A_2 N] \not\equiv A_2$
 - $[A_{int} A_2 N] \equiv N$

Such an inference pattern is subject to contextual variables, many of which are not available to sentential compositional mechanisms, but some constraints can be identified. For example, the closer the head noun is to a sortal base level category, such as *bird*, *table*, or *tool*, the more likely the inference in (10b) will go through. Consider the examples below:

- (11) a. The store bought an alleged antique vase.
 b. The researcher found an alleged Mozart sonata.

These cases make it clear that the epistemic uncertainty in (11) involves an additional aspect of the NP, beyond the unassailable characteristics of the entailed head. That is, the object is clearly a vase (in (11a)) and demonstrably a sonata (in (10b)). Such evidence, however, will not always be available within the composition of a sentence, but will be derivable from context (if at all). We will refer to the canonical inference in (10a) as the “Wide-scope reading”, and the inferences in (10b-c) as the “Narrow-scope reading”.

Another interesting distinction emerging in the basic $[A N]$ construction with intensional adjectives is one based on the type of the nominal head. The most common semantic types occurring in the corpus are shown below, along with apparent scoping behavior.

- (12) a. EVENT NOMINAL: *violation*, *misconduct*, *murder*, *assault*. The more specific nominal descriptions carry greater inferential force for the hypernym. That is, *murder* suggests inference of a death.
 b. AGENTIVE NOUN: *collaborator*, *perpetrator*, *murderer*, *criminal*. Epistemic scope is over the entire sortal. The canonical form, “the alleged criminal”.

Certainly more is intended than a mere hypernymic assertion, including the associated presuppositions of the context variables introducing the allegation and the epistemic uncertainty itself. These are issues presently being explored.

- c. UNDERGOER NOUN: *victim*. While not always the case, the scope is narrowed to a modification of the event: For example, “the alleged victims of Whitey Bulger”.

Consider the sentences in (13), where *alleged* is modifying an event nominal.

- (13) a. He denies the alleged assault on the police.
 b. The greatest number of alleged violations occurred in California.
 c. He’s been charged in connection with the alleged murder of John Smith, whose mutilated body ...

The inferences associated with (13a-b) follow from the template in (10a). For sentence (13c), however, we need to infer that there was, in fact, a killing, although it is uncertain whether it was a murder. This requires the inference rule below, where the hypernym of the event nominal is inferable from the context.

- (14) Given the construction $[A_{int} N]$, where N is an event nominal, with certain feature, then:
- $[A_{int} N] \not\equiv N$
 $\equiv N'$ where $N \subseteq N'$

We refer to this inference rule as the “Hypernym reading”. Similar remarks hold for undergoer nominals in some contexts, where the scope of the intensional adjective can be lowered to a modification of the event description. This is illustrated below, in (15b).

- (15) a. Testimony will be heard from the alleged victim in court.
 b. The families of two alleged victims of James “Whitey” Bulger have received compensation.

Sentence (15a) behaves according to the canonical template, while (15b) involves a narrower scope of the epistemic uncertainty. That is, the inference should be made that there are victims, but the cause (or etiology) of this designation is uncertain. This rule is formally related to that presented above in (10), where the modification (argument specification, in fact) is postnominal.

- (16) Given the construction $[A_{int} N XP_{mod}]$, where XP_{mod} is a modification or argument, then:
- $[A_{int} N XP_{mod}] \not\equiv N XP_{mod}$
 - $[A_{int} N XP_{mod}] \equiv N$

Summarizing the semantic behavior for this class, we have identified at least three distinct structure-to-inference mappings associated with intensional (non-subjective) adjectives. These are:

(17) Structure-to-Inference Mappings:

a. Wide-scope reading:

$[A_{int} N] \not\models N$

b. Narrow-scope reading 1:

$[A_{int} A_2 N] \not\models A_2, \models N$

c. Narrow-scope reading 2:

$[A_{int} N X P_{mod}] \models N$

d. Hypernym reading:

$[A_{int} N] \models N'$ where $N \subseteq N'$

3 Data Collection and Discussion

There are approximately 50 intensional (sub-selective) adjectives that we have identified, from which we have selected the most frequent 30 for our investigation. Fewer than 10 of these are root adjectives (*superficial*, *putative*), and most are participial adjectival derivations, such as *alleged*, *supposed*, and *believed*. For each adjective, we have extracted 100 snippets from the corpus, where snippets are three-sentence fragments from the text. This gives us a corpus of 3,000 snippets for intensional adjectives.

We are developing an initial classification of 1,000 of these adjectives based on the inferential patterns discussed in the previous section; i.e., wide-scope, narrow-scope, and hypernym readings. These are the initial structure-to-inference templates which will constitute the small gold standard. This annotation is being performed by undergraduate linguistics majors, with three annotations per snippet. That is, we construct the examples that fit the identified test patterns, as shown in (18) and (19) below. In these examples, the inference in (18) is legitimate, while that in (19) is false.

(18) Hypernym Reading:

(T): A teenage girl has been arrested over the **alleged murder** of a mourner at a funeral in London.

(H): A mourner died.

(19) Wide-Scope Reading:

(T): She was then tried and executed in 1952 by Stalin as an **alleged spy**.

(H): She was a spy.

We then will submit these stimuli to MTurkers with the same guidelines as those given to the linguists, and examine the differences in judgments. That is, for those cases that do not accord with the pre-assigned classification, we try to isolate the factors contributing to when the judgment goes against the expected inference. To this end, we perform a statistical analysis of the contexts of the adjective for both the cases that are in accordance with the classification and the cases that are not.

The goal of this ongoing effort is to elucidate the semantic properties and inferential patterns associated with adjectives in natural language. As we have tried to make clear from this brief report, the semantic behavior of adjectives in actual language use are much more nuanced and subtle than previously documented. We hope to report on further results and insights in the near future.⁴

Acknowledgements

This research was supported by a grant from the NSF (NSF-IIS 1017765). I would like to thank Zachary Yochum, Annie Zaenen, and Christiane Fellbaum for their comments and discussion. I would also like to acknowledge the anonymous reviewers for the workshop for their helpful comments. All errors and mistakes are, of course, my own.

References

- M. Amoia and C. Gardent. 2006. Adjective based inference. In *Proceedings of the Workshop KRAQ'06 on Knowledge and Reasoning for Language Processing*, pages 20–27. Association for Computational Linguistics.
- M. Amoia, C. Gardent, et al. 2008. A test suite for inference involving adjectives. *Proceedings of LREC'08*, pages 19–27.
- Gemma Boleda, Marco Baroni, Louise McNally, and Nghia The Pham. 2013. Intensionality was only alleged: On adjective-noun composition in distributional semantics. In *Proceedings of IWCS 2013*.
- R.M.W. Dixon. 1991. *A new approach to English grammar on semantic principles*. Oxford University Press.

⁴A related paper, (Boleda et al., 2013), on the semantics of intensional adjectives, is being presented at the same venue as the present paper, and came to my attention only recently. As a result, the analysis therein has not been referenced in this paper.

- H. Kamp and B. Partee. 1995. Prototype theory and compositionality. *Cognition*, pages 57–129.
- H. Kamp. 1975. Two theories about adjectives. In *Formal Semantics of Natural Language*, pages 123–155. University Press.
- John Lyons. 1977. *Semantics*. Cambridge University Press.
- V. Raskin and S. Nirenburg. 1995. Lexical semantics of adjectives. *New Mexico State University, Computing Research Laboratory Technical Report, MCCS-95-288*.