

So-Called Non-Subsective Adjectives

Ellie Pavlick and Chris Callison-Burch
University of Pennsylvania

September 21, 2016

Abstract

The interpretation of adjective-noun pairs plays a crucial role in tasks such as recognizing textual entailment. Formal semantics often places adjectives into a taxonomy which should dictate adjectives entailment behavior when placed in adjective-noun compounds. However, we show experimentally that the behavior of subsective adjectives (e.g. *red*) versus non-subsective adjectives (e.g. *fake*) is not as cut and dry as often assumed. For example, inferences are not always symmetric: while *ID* is generally considered to be mutually exclusive with *fake ID*, *fake ID* is considered to entail *ID*. We discuss the implications of these findings for automated natural language understanding.