

Adjectival Modification in Text Meaning Representation

Victor Raskin and Sergei Nirenburg

Computing Research Laboratory

New Mexico State University

Las Cruces, N.M. 88003, U.S.A.

raskin, sergei@crl.nmsu.edu

Abstract

This work belongs to a family of research efforts, called microtheories and aimed at describing the static meaning of all lexical categories in several languages in the framework of the MikroKosmos project on computational semantics. The latter also involves other static microtheories describing world knowledge and syntax-semantics mapping as well as dynamic microtheories connected with the actual process of text analysis. This paper describes our approach to detecting and recording adjectival meaning, compares it with the body of knowledge on adjectives in literature and presents a detailed, practically tested methodology for the acquisition of lexical entries for adjectives. The work was based on the set of over 6,000 English and about 1,500 Spanish adjectives obtained from task-oriented corpora.

1. The Ontological Approach

The work on adjectives reported in this paper constitutes a descriptive “microtheory” in the MikroKosmos semantic analyzer (Onyshkevych and Nirenburg 1994; and Beale et al. 1995), designed to serve as a component of a knowledge-based machine translation system (Nirenburg et al. 1992).

MikroKosmos combines findings from a variety of quasi-autonomous microtheories of language phenomena, world knowledge organization and procedural knowledge at the level of computer system architecture. The basic motivation for this organization is the continued inability of the fields of linguistics and NLP to produce a general-coverage, unified theory of treatment of language phenomena, a failure especially pronounced in areas beyond computational syntax.

The purpose and result of the MikroKosmos analysis process is the derivation of an interlingual representation for natural language inputs. The language in which these representations are expressed is called the “text meaning representation” (TMR) language, and “texts” in this language are called, simply, TMRs. TMR is a frame-based language, where frame names typically refer to instances of ontological concepts, slot names are derived

from a set of ontological properties and slot fillers are either elements of property value sets or pointers to concept instances.

An ontology, a world model containing information about types of things, events and properties in the world, is a necessary prerequisite for a TMR language. “An ontology for NLP purposes is a body of knowledge about the world (or a domain) that a) is a repository of primitive symbols used in meaning representation; b) organizes these symbols in a tangled subsumption hierarchy; and c) further interconnects these symbols using a rich system of semantic and discourse-pragmatic relations defined among the concepts” (Mahesh and Nirenburg 1995: 1). The function of the ontology is to supply “world knowledge to lexical, syntactic, and semantic processes” (ibid).


The lexicon in MikroKosmos “mediates between the TMR and ontology” (Onyshkevych and Nirenburg 1994: 2). Lexicon entries for most open-class lexical items represent word and phrase senses, which can be either directly mapped into ontological concepts or derived by locally (that is, in the lexicon entry itself) modifying constraints on property values of concepts used to specify the meaning of the given lexical item. Lexical-semantic information as well as clues for contextual semantic and pragmatic processing are typically located in the lexicon, adjectives being no exception. In the following section we illustrate the structure of those parts of the lexicon entry in MikroKosmos which bear on the description of adjectival meaning.

2. The Ontological Approach to the Meaning of a Typical Adjective

A simple, prototypical case of adjectival modification is a scalar adjective, which modifies a noun both syntactically and semantically. Our microtheory associates its meaning with a region on a scale which is defined as the range of an ontological property (cf. Carlson and Nirenburg, 1990). The contribution that the adjective makes to the construction of a semantic dependency structure

(TMR) typically consists of inserting its meaning (a property-value pair) as a slot filler in a frame representing the meaning of the noun which this adjective syntactically modifies.

Thus, in *big house*, *big* will assign a high value as the filler of the property slot SIZE of the frame for the meaning of *house*. The range of the ontological property SIZE is a numerical and continuous scale. Each numerical



Big will, then, be assigned a value of '> 0.75' value on the SIZE scale. These values are a crucial part of the lexical mapping (LEX-MAP) from language units to TMR units included in the semantics (SEM-STRUC) "zone" of their lexical entries. Equally crucial is the syntactic-semantic dependency mapping (linking) between the syntactic-structure (SYN-STRUC) and SEM-STRUC zones, which in MikroKosmos is carried out with the help of special variables. The syntactico-semantic information in the lexicon entry for *big* is as follows:¹

```
(1) (big
      (big-Adj1      ;the first adjectival sense of BIG
      (CAT adj)
      (SYN-STRUC    ;syntactic characteristics:
      (1 ((root $var1) ;subcategorization pattern
          (cat n)      ;1 (attributive); $var1 is
          (mods ((root $var0)))) ;bound to the noun
          ;the adjective modifies;
          ;$var0 is bound to the ad-
          ;jective itself
      (2 ((root $var0) ;subcategorization pattern
          (cat adj)    ;2 (predicative)
          (subj ((root $var1) ;this standard Adj
                ;SYN-STRUC is
                (cat n)))))) ;omitted from the other
          ;examples
      (SEM-STRUC
      (LEX-MAP      ;the syntax-semantics map-
      ((1 2) (size-attribute ;ping valid for both
              (domain (value ^$var1);patterns; '^'
                       ;means "the meaning of"
              (sem physical-object)) ;selectional
              ;restriction
              (range (value (> 0.75)) ; the value is
                       ;in the top 25 percentile
                       ;of the scale
              (relaxable-to (value (>0.6)))))))));re-
          ;laxed values are for pro-
          ;cessing metonymies
```

The standard procedure for representing adjectival modification in TMRs is, then, to insert the scale name

¹ Many zones which are actually present in the entries for these adjectives in the MikroKosmos lexicon are omitted from the examples.

scale can be measured in an absolute manner (e.g., LINEAR-SIZE in feet, yards, or millimeters, or TIME in seconds). But often natural language expressions do not refer to absolute magnitudes but rather to abstract relative ones, as in the case of *big*. We assume a 0 to 1 numerical range for such abstract scales. For abstract references to SIZE, the fillers in English can be:

and scale value for an adjective as a property-value pair in the frame describing the meaning of the noun the adjective modifies. For a noun like *house*, whose appropriate sense (2) is directly mapped into an ontological concept, the meaning of *big house* will be represented as a TMR fragment shown in (3):

```
(2) (house
      (house-N1
      (CAT n
      (SYN-STRUC
      (1 ((root $var0)
          (cat n))))
      (SEM-STRUC
      (LEX-MAP
      (2 (private-home)
          (size-attribute (value > 0.75))
```

More complex cases of adjectival modification are discussed in Section 4.

3. Semantic and Computational Treatment of Adjectives: Old and New Trends

The literature on adjectives shows a scarcity of systematic semantic analyses or lexicographic descriptions of adjectives. Most of the linguistic scholarship focuses on the taxonomies of adjectives, on the differences between the attributive and predicative syntactic usages as well as other syntactic transformations associated with various adjectival usages, on the qualitative/relative distinctions among adjectives, which is related to the predicative/attributive usages, and on the gradability/comparability of qualitative adjectives (for a detailed survey, see Raskin and Nirenburg 1995: 3-20).

As computational semantics moves to large-scale systems serving non-toy domains, the need for large lexicons with entries of all lexical categories is becoming increasingly acute, and the attention is turning more towards such previously neglected or avoided categories as the adjectives. Recently, there have appeared some first indications of this attention—see, for instance,

Smadja (1991), Beckwith *et al.* (1991), Bouillon and Viegas (1994), and Pustejovsky (1995). This research is a step in the same direction.

Our analysis of adjectives, with the goal of supporting semantic analysis, shows that the issues important for adjective meaning representation are quite different from those debated in literature on adjectives. Thus, it becomes clear that the *scalar/non-scalar* dichotomy, and not the *attributive/predicative* distinction which dominates the literature, is the single most important distinction in semantic treatment of adjectives. The continuous numerical scales associated with the true scalars also render the issue of gradability and comparability rather trivial (see Raskin and Nirenburg 1995: 25-26).

Another essential issue is the grain size of description. In (3) the linking attribute (SIZE) is selected rather high in the hierarchy of attributes, because in the ontology SIZE-ATTRIBUTE is the parent of such properties as LENGTH-ATTRIBUTE, WIDTH-ATTRIBUTE, AREA-ATTRIBUTE, WEIGHT-ATTRIBUTE, etc. If the context does not allow the analyzer to select one of those, a coarser-grain solution is preferred. In other words, we represent the meaning of *big house* without specifying whether *big* pertains to the length, width, height or area of a house. This is the result of a principled decision, based on the *principle of practical effability*² (Raskin and Nirenburg 1995: 46ff), which stipulates that, in MT, the target language should be expected to have a corresponding adjective of a comparably large grain-size.

This issue has been often discussed on the example of the adjective *good* (cf. Katz 1972, Pustejovsky 1995). We deliberately settle on a grain size of description coarser than the most detailed semantic analysis possible (4).

```
(4) (good
      (good-Adj1
       (CAT adj)
       (SYN-STRUC
        (1 ((root $var1)
            (cat n)
            (mods ((root $var0))))))
        (2 ((root $var0)
            (cat adj)
            (subj ((root $var1)
                  (cat n))))))
       (SEM-STRUC
        (LEX-MAP
         (attitude
          (type evaluative)
          (attitude-value (value (>0.75))
                           (relaxable-to (value (>0.6))))))
```

² Derived from Tarski's and Katz's effability principle (Tarski 1956: 19-21; Katz 1978: 209) and extended to NLP.

```
(scope ^$var1)
(attributed-to *speaker*)))))
```

The finest grain-size analysis requires that a certain salient property of the modified noun is contextually selected as the one on which the meaning of the noun and that of the adjective is connected. In our approach, the representation solution for *good* would be to introduce an evaluation attitude, with a high value and scoped over this property. Salient properties are, however, hard to identify formally, as is well known, for instance, in the scholarship on metaphor, where salience is the determining factor for the similarity dimension on which metaphors (and similes) are based. It is, therefore, wise to avoid having to search for the salient property, and the principle of practical effability offers a justification for this.

4. Non-Property-Based Adjectival Modification

This section contains a brief discussion of the semantic treatment of adjectives which cannot be reduced to the standard property-based type of adjectival modification. This discussion illustrates an important point in our approach, namely, that syntactic modification does not necessarily imply semantic modification.

4.1 Attitudes

Good is, of course, a scalar. Nevertheless, unlike in the case of *big* (2), the LEX-MAP for (4) does not contain a property-value pair that can be attached to the frame of the modified noun like *house* in the TMR. Instead, the meaning representation of *good* introduces an **attitude** on the part of the speaker with regard to the modified noun. In the TMR, the attitudes characterize the whole proposition, and thus the semantic link between the modified noun and the adjective is weakened. There are other types of adjectives which challenge the common-sense view that the meaning of the adjective somehow "amalgamates" with the meaning of the modified noun, and most of these types are non-scalar or only marginally scalar.

4.2 Temporal Adjectives

The purely temporal knowledge in MikroKosmos is recorded with the meaning of the entire proposition, and adjective entries are not marked for it. Some temporal adjectives, of the kind that Levi presents as derived from adverbs rather than nouns (examples (1.9) in Levi 1978: 7, repeated here as 5), are analyzed in a different manner

precisely because they do not modify semantically the nouns they modify syntactically--in other words, the temporal meaning of the adjective characterizes the proposition. Thus, *occasional visitor* (5iii) is analyzed as a rhetorical paraphrase of *visit occasionally*.

- (5) (i) former roommate
 (ii) early riser
 (iii) occasional visitor
 (iv) eventual compromise

4.3 Membership adjectives

The membership class has been largely ignored in the literature. There has been a sporadic interest in the adjective *fake* (see Iwanska 1995--cf. Raskin 1981) because it clearly violates the simplistic subset-forming notion of adjective meaning, such that red houses are a subset of all houses. But there are many other adjectives which use exactly the same type of lexical entry, and their similarity to each other and to *fake* had not been noticed before.

The most typical adjectives in the member subclass are *authentic* (6), *fake* (7), and *nominal* (8). Many others are their synonyms and near-synonyms. The lexical entry for this subclass focuses on two major elements: first, whether the modified noun is a member of a certain set--authentic and nominal members are but fake members are not; and, second, whether the properties of this noun intersect significantly with those of the set members--the properties of authentic members overlap with the common properties of the set members on most important properties; the properties of fake members overlap with those of the set members only on unimportant properties, such as physical resemblance--e.g. *fake gun*; and the properties of nominal members overlap more significantly with those of the set members but not on the most important ones.

The first element is represented in a set notation: set1 shows that ^\$var1 belongs to the set, whose typical member is denoted by a variable refsem1, in the case of *authentic* and *nominal* but not in the case of *fake*. Set2 is the set of all properties of the members of set1; set3 is the set of all properties of ^\$var1; set4 is, essentially, the intersection of set2 and set3.

The second element is represented as the value of a saliency (importance) attitude to the intersection between the properties of the modified noun and those of the set members it is purported to belong to: the saliency value is 1.0 for *authentic*, still high for *nominal*, and low for *fake*. This representation is based on the assumption that functioning as a member, which differentiates between *authentic* and *nominal*, in that the former does and the latter does not function as a member should, is the most salient feature, while something

like physical similarity (a fake gun only looks like a gun) is the least salient one. (7) and (8) below are shown only partially, where they contrast with (6).

- (6) (authentic
 (authentic-Adj1
 CAT adj)
 (SEM-STRUC
 (LEX-MAP
 ((1 2) (set1
 (member refsem1) ;refsem X are vari-
 (member ^\$var1)) ;ables not used by
 (set2 ;the linking process
 (member refsem1.*)) ;"*" means all
 (set3 ;properties in a
 (member ^\$var1.*)) ;a concept
 (set 4
 (member (AND (set2.member
 set3.member))))
 (attitude
 (type salience)
 (attitude-value 1.0)
 (scope set4
 (attributed-to *speaker *))))))
- (7) (fake
 (fake-Adj1
 ((1 2)(set1
 (member refsem1)
 ((member ^\$var1)
 (polarity negative)))
 (attitude-value (value (< 0.25))))
- (8) (nominal
 (nominal-Adj1
 (attitude-value (value (< 0.75))))

4.4 Event-Related Adjectives

To derive the semantic part of an adjectival entry from a verbal entry, first one must identify the case, or thematic role (such as agent, theme, beneficiary, etc.) filled by the noun modified by the adjective in question. We illustrate this process using the lexical entries for *abusive* and *abuse*. The superentry for *abuse* includes at least three senses, roughly, *abuse-V1* 'insult verbally,' *abuse-V2* 'violate a law or a privilege,' and *abuse-V3* 'assault physically,' and the adjective may be derived from any one of them. What is *abusive* is either the event (E) itself, as in *abusive speech* or *abusive behavior*, or the agent (A) of the event, as in *abusive man* or *abusive neighbor*. *Abusive_{IE}* is then the eventive sense of the adjective formed from *abuse-V1* (9), and *abusive_{IA}* is the agentive sense of the adjective in the same sense of *abuse*. The difference between the two is, essentially, in the position of ^\$var1 in the LEX-MAP and in the scope of attribution of the two attitudes inherited from the verbal entry. Naturally, the adjective entries replace the verbal SYN-STRUC below with the standard Adj one (see

(1) above--for more data and discussion see also Raskin and Nirenburg 1996).

```
(9) (abuse
    (abuse-V1
     (CAT V)
     (SYN-STRUC
      ((root $var0)
       (cat v)
       (subj ((root $var1)
              (cat n))
              (obj ((root $var2)
                    (cat n))))))
     (SEM-STRUC
      (LEX-MAP
       (communicative-event
        (agent (value ^$var1)
              (sem human))
        (benef (value ^$var2)
              (sem human))
        (theme (value refsem1))
        (attitude1
         (type evaluative)
         (attitude-value (value (< 0.25)))
         (scope refsem1)
         (attributed-to (OR (^$var2 speaker))))
        (attitude2
         (type evaluative)
         (attitude-value (value (< 0.25)))
         (scope ^$var2)
         (attributed-to ^$var1))))))
```

4.5 Relative (Denominal) Adjectives

Relative adjectives are denominal, object-related, in their meaning. The following example illustrates the connection between nominal and adjectival meanings.

```
(10) (i) (medicine
        (medicine-N1)
        (CAT n)
        (SYN-STRUC
         (root $var0)
         (cat n)))
      (SEM-STRUC
       (LEX-MAP
        medicine))))
    (ii) (medical
          (medical-Adj)
          (CATadj)
          (SEM-STRUC
           (LEX-MAP
            (^$var1
             (pertain-to medicine))))))
```

As the default property connecting the modifier to the modified, the MikroKosmos analyzer uses the catch-all relation PERTAIN-TO. We have identified several more specific relations.

The first such relation is OWNED-BY, as in *federal-Adj1* in the sense of “owned by a federation.” Another

specific relation is HAS-AS-PART, as in *malignant-Adj3* in the sense of containing cancer cells. LOCATION is also a common relation, as in *international-Adj1*, “taking place in a set of two or more countries.” It is interesting that another sense of *international* utilizes the OWNED-BY property noted above, as in “owned by a set of two or more countries,” and yet another combines LOCATION with event-relatedness, as in “manufactured in a set of two or more countries.”

The disambiguation among such multiple senses is not a simple matter, and in an unusual contraposition to the standard semantic problem of infinite polysemy, a move up, rather than down, to the undifferentiated generic meaning of an adjective like *international* is recommended in case of disambiguation problems. In other words, while we continue to discover more specific relations between the lexical entries of denominal adjectives and the nouns they are derived from, the generic PERTAIN-TO property should not be discarded. This move is, again, related to the issue of grain-size of semantic description.

5. Adjectives and Other Modifiers

The MikroKosmos analyzer treats modification by attempting to merge the meanings of the modifiers into the meanings of the modified. For those modifiers whose meanings are (possibly, sets of) property-value pairs, the method is to insert the values they carry into the same property slot in the modified. For instance, the sense of *smooth* as in *smooth silk* will be a range on the TEXTURE scale. If TEXTURE is defined as a property of PHYSICAL-OBJECT or MATERIAL, and SILK is a descendent of either of them, then the value carried in the lexicon entry for *smooth* will be inserted by the analyzer as the TEXTURE property value for the instance of *silk* in the TMR.

Our approach covers all property modification in language, not only adjective-noun combinations. Thus, it would be applicable to noun-noun combinations, adverb-verb combinations and other modification situations, as illustrated in (11):

(11) Modified	Modifiers
Verb	Adverb, Noun, Prepositional Phrase
Noun	Adjective, Prepositional Phrase
Adjective	Adverb, Prepositional Phrase
Adverb	Adverb

The most challenging cases in all kinds of modification would be those where syntactic dependency does not predetermine semantic dependency. In this paper we have illustrated a method, based on ontology and text meaning representation, of treating such discrepancies in dependency for adjectival modification. This method

has been tested in the MikroKosmos semantic analyzer based on the lexical entries for 6,000 Spanish and 1,500 English adjectives.

The method is based on the discovery of a small number of basic types of adjectival lexical entries and its use, with minor modifications, with a large number of specific lexical entries, thus making the acquisition of adjectives cognitively easier, faster, and cheaper. Each type of lexical entry determines a type of modification relationship between the adjective and the kind of nouns it modifies, most significantly, whether this relationship is property-based or not-property-based. We have also discovered that this approach to adjectival meaning is language-independent: what varies from language to language is the adjectival superentries, i.e., the various combinations of different meanings of the same adjective, as well as adjectival availability for a certain meaning. i.e., whether a specific meaning can be expressed adjectivally in a language. Most adjectival meanings of one language are, however, expressed adjectivally as well in another language, and the lexical entry for this meaning is then unchanged.

In many languages, adjectives and adverbs are the same. Is our approach to adjectival modification of nouns applicable to adverbial modification of verbs? Initial research shows that the property-/non-property-based dichotomy holds there as well. We intend to test the hypothesis that this method extrapolates to all the above types of modification as well.

Acknowledgments

The research reported in this paper was supported by Contract MDA904-92-C-5189 with the U.S. Department of Defense. Victor Raskin is grateful to Purdue University for permitting him to consult CRL/NMSU on the MikroKosmos project. Both authors feel indebted to the other members of the MikroKosmos team.

References

- Beale, Stephen, Sergei Nirenburg, and Kavi Mahesh 1995. Semantic Analysis in the MikroKosmos Machine Translation Project. In **Proceedings of the Second Symposium on Natural Language Processing (SNLP-95)**, August 2-4. Bangkok, Thailand.
- Beckwith, Richard, Christiane Fellbaum, Derek Gross, and George A. Miller 1991. WordNet: A lexical database organized on psycholinguistic principles. In: Uri Zernik (ed.), **Lexical Acquisition: Exploiting On-line Resources to Build a Lexicon**. Hillsdale, N.J.: Erlbaum, pp. 211-232.
- Bouillon, Pierrette, and Evelyne Viegas 1994. A semi-polymorphic approach to the interpretation of adjectival constructions: A cross-linguistic perspective. In: **Proceedings of the Sixth EU-RALEX International Congress**, Amsterdam: Free University of Amsterdam, pp. 36-44.
- Carlson, Lynn, and Sergei Nirenburg 1990. World Modeling for NLP. Technical Report CMU-CMT-90-121, Center for Machine Translation, Carnegie Mellon University, Pittsburgh, PA.
- Iwańska, Lucja 1995. Semantics, pragmatics, and context of intensional negative adjectives: Not a toy problem. Paper presented at the Context in Natural Language Processing. IJCAI '95 Workshop, Montreal, August 1995.
- Katz, Jerrold J. 1972. Semantic theory and the meaning of good. *Journal of Philosophy* 61, pp. 736-760.
- Katz, Jerrold J. 1978. Effability and translation. In: F. Guenther and M. Guenther-Reutter (eds.), **Meaning and Translation: Philosophical and Linguistic Approaches**. London: Duckworth, pp. 191-234.
- Levi, Judith N. 1978. **The Syntax and Semantics of Complex Nominals**. New York: Academic Press.
- Mahesh, Kavi, and Sergei Nirenburg 1995. A situated ontology for practical NLP. A paper presented at the IJCAI '95 Workshop on Basic Ontological Issues in Knowledge Sharing. Montreal, August 19-21.
- Nirenburg, Sergei, Jaime Carbonell, Masaru Tomita, and Kenneth Goodman 1992. **Machine Translation: A Knowledge-Based Approach**. San Mateo, CA: Morgan Kaufmann.
- Onyshkevych, Boyan, and Sergei Nirenburg 1994. The Lexicon in the Scheme of KBMT Things. Memoranda in Computer and Cognitive Science MCCS-94-277. Las Cruces, N.M.: New Mexico State University.
- Pustejovsky, James 1995. **The Generative Lexicon**. Cambridge, MA: MIT Press.
- Raskin, Victor 1981. How to handle fake guns. Meeting Handbook, Annual Meeting of the Linguistics Society of America, New York.
- Raskin, Victor, and Sergei Nirenburg 1995. Lexical Semantics of Adjectives: A Microtheory of Adjectival Meaning. Memoranda in Computer and Cognitive Science MCCS-95-288. Las Cruces, N.M.: New Mexico State University.
- Raskin, Victor, and Sergei Nirenburg 1996. Lexical rules for deverbal adjectives. Paper presented at the ACL '96 Workshop on the Breadth and Depth of Semantic Lexicons. Santa Cruz, CA.
- Smadja, Frank 1991. Macrocoding the lexicon with co-occurrence knowledge. In: Uri Zernik (ed.), **Lexical Acquisition: Exploiting On-line Resources to Build a Lexicon**. Hillsdale, N.J.: Erlbaum, pp. 165-189.