"Tall'GoodV'High" — Compared to What?

Steffen Staab & Udo Hahn

Computational Linguistics Research Group
Freiburg University, Werthmannplatz 1, D-79085 Freiburg, Germany
{staab,hahn}@coling.uni-freiburg.de
http://www.coling.uni-freiburg.de

Abstract

We specify a model for the conceptual interpretation of positive gradable adjectives. Building on a classification-based terminological reasoning approach we define comparison classes and class norms and specify how a degree is related to its corresponding class norm.

1 Introduction

There is a wide-spread consensus in the linguistic community that the semantics of positive adjectives like "rail" is captured by a binary predicate relating a degree to a comparison class (cf., e.g., [Klein, 1980; Bierwisch, 1989]). Hence, *"Peter* in (Ib) should not be referred to as *"tall"* in a general sense, but as *"tall in comparison to a class* C", where C is constrained by the context in which *"tall"* occurs. This becomes evident considering the full version of example (1), where the utterance (Ia) crucially determines the valid comparison class for *"tall(Pcter,C)'* C being the class of "4-year-old-boys".

(1) a. *Peter is 4 years old.* b. *Peter is tall.*

From a computational point of view, two questions arise. First, how are degree expressions - positive gradable adjectives, in particular - represented? And, second, how are valid comparison classes for positive adjectives determined? The representation problem has already been tackled from a quantitative perspective employing interval representations [Simmons, 1993] and fuzzy logic [Zadeh, 1978], but also qualitative approaches have been considered [Kamei and Muraki, 1994; Schwartz, 1989]. While both schools have notorious difficulties in incorporating the corresponding complementary methodology, in addition, they both fail to treat ordinal information appropriately such as for transitive reasoning on comparative constructions. As an alternative, we propose a qualitative representation and commonsense inferencing scheme that overcomes these deficits (cf. Section 2).

Until now, only [Bierwisch, 1989] has touched on the problems how to determine comparison classes, though in a rather sketchy way. Hence, our algorithm in Section 3 constitutes the first attempt at providing an explicit computation procedure for the conceptual interpretation of positive gradable adjectives in terms of finding its proper comparison class.

2 The Degree Calculus

A general representation and inferencing scheme for natural language degree expressions must account for several requirements simultaneously. First, it must allow for different forms of degree expressions, esp. comparatives and positive adjectives. Second, it must allow for a basically qualitative representation, since many expressions cannot be attributed any numbers at all, or only in an *ad hoc* way. Third, it must allow for transitive inferences, since these are the most basic reasoning patterns for degrees. Fourth, at least for technical domains, measure phrases must be incorporated, too. None of the approaches mentioned in Section 1 are able to cope with these criteria, while the *Degree Calculus* we sketch below offers a synthesis of these four requirements.

2.1 Degree Representation

Following a widely acknowledged proposal, we describe the meaning of positive gradable adjectives (as in (2a)) by a comparison to a norm degree, N_C, of a comparison class C (cf. (2b), where *H* denotes the function *Height*).

(2) a. *Peter is tall.*
b. $H(Peter) \succ_{\Delta(0)} N_C$.

In order to account for degree modifiers, too, qualitative labels are introduced that may modify such comparisons. These *modifier labels* are (partially) ordered according to their "strength" (cf. (3)).

(3) $\Delta(much) > \Delta(slightly) > \Delta(0) \ \wedge$
$\Delta(very) > \Delta(slightly) > \Delta(0) \ \wedge \ \ldots$

$\Delta(0)$ renders the comparison relation "$\succ_{\Delta(0)}$" roughly equivalent to the common relation "$>$". But other labels modifying the comparison relation require *more than a certain distance* (given by the label, e.g., $\Delta(very)$) between the respective degrees to render the comparison valid (cf. (4b)).

(4) a. *Peter is very tall.*
b. $H(Peter) \succ_{\Delta(very)} N_C$.

Furthermore, a negative label as in (5b), technically speaking, leaves open the interpretation whether Mary is taller than Peter or Peter is *less than* Δ(slightly) taller than Mary. Since "slightly taller" describes a twofold restriction — Peter being taller than Mary but not that much — the expression (5c) must be conjoined with (5b) to properly denote (5a).

(5)
a. *Peter is slightly taller than Mary.*
b. $H(Mary) \succ_{-\Delta(slightly)} H(Peter) \wedge$
c. $H(Peter) \succ_{\Delta(0)} H(Mary)$.

The approach we have just sketched provides a valid representation framework for positives, comparatives, equatives, and superlatives (as all-quantified comparatives). In addition, it may easily account for measure phrases as in (6) by augmenting the ordering on the modifier labels (3) in terms of scale-specific ordering relations for measure phrase modifiers, such as $\Delta(5 \text{ inches}) > \Delta(\text{much})$.

(6)
a. *Peter is more than 5 inches taller than Paul.*
b. $H(Peter) \succ_{\Delta(5\,inches)} H(Paul)$.

2.2 Rules of Inference for Degrees

As far as measure phrase modifiers are concerned (cf. example (6)), a calculus can be supplied the axioms of which are stated in Table 1.[1]

	$\forall a, b, c \in D$, the set of degree instances, and $\forall x, y \in L$, the set of modifier labels:	
1.	$x < \Delta(0) \Rightarrow (a \succ_x a)$	(reflexivity)
2.	$x \geq \Delta(0) \Rightarrow (a \succ_x a \Rightarrow \bot)$	(contradiction)
3.	$(a \succ_x b \wedge a \succ_y b) \Leftrightarrow a \succ_{\max(x,y)} b$	(subsumption)
4.	$(a \succ_x b \wedge b \succ_y c) \Rightarrow a \succ_{x+y} c$	(composition)

Table 1: Selected Axioms of the Degree Calculus

However, we may also adapt this calculus to inferencing with qualitative labels. While the axioms 1–3 can be applied directly, the composition axiom can be used to infer *minimal assertions*, at least. As an example, consider the inference step in (7), the formal description of which is given in (8). Considering the axioms only, no value can be computed for $\Delta(\text{very}) + (-\Delta(\text{slightly}))$. Including the ordering in (3), however, allows one to conclude that this value subsumes the less restrictive $\succ_{\Delta(0)}$.

(7)
Peter is slightly taller than Mary.	*given*
Peter is very tall.	*given*
Mary is tall.	*inferred*

(8)
$$\frac{\dfrac{H(Mary) \succ_{-\Delta(slightly)} H(Peter) \wedge \ldots \quad H(Peter) \succ_{\Delta(very)} N_C}{H(Mary) \succ_{\Delta(very)-\Delta(slightly)} N_C}\ (composition)}{H(Mary) \succ_{\Delta(0)} N_C}\ (subsumption)$$

This way of reasoning already incorporates simpler types of transitive reasoning, like "Peter is taller than Paul and Paul is taller than John." implies "Peter is taller than John." Nevertheless, the addition of scale-specific constraints, such as $\Delta(\text{slightly}) < \Delta(2 \text{ inches})$ on the scale for persons' heights, can still further augment the system's inference capabilities.

[1]In our complete axiomatization, we also use the operator "$\succeq$". Extensions like "$\prec_x$", "equal with tolerance x", etc. are often useful and convenient. They can be easily defined with the help of "$\succ_x$" and "$\succeq_x$", cf. [Staab and Hahn, 1997b].

3 Comparison Classes

In this section, we specify an algorithm that relates the degree described by a positive gradable adjective a to a *class norm.*[2] This class norm is a degree of the same type (e.g., HEIGHT) as the one described by the adjective (e.g., *"tall").* The class norm belongs to a *comparison class,* (e.g., the set of 4-year-old-boys) which is a set of individuals or - in terminological terms - a concept *C* with instances o_i. If the degree of such an instance o_i of *C* exceeds the class norm, then it can be asserted that "o_i is *a* for C". At the core of the algorithm lies *knowledge of interrelations* which must be available for inferencing and will subsequently be considered in more detail.

3.1 Representation of Comparison Classes

Sentence (9) contains an occurrence of a positive gradable adjective. Computing its comparison class is common part of the determination of the conceptual interpretation of that utterance, since it is explicitly given by the underlined phrase.

(9) *The XI1 offers very good quality for* ***a laser printer that costs $800.***

Considering this example, we focus on how the declarative representation of a comparison class is *dynamically* created from the utterance and the concepts given in a domain knowledge base. The terminological system we use (cf. [Woods and Schmolze, 1992] for a survey) allows to define a comparison class COMP-CLASS-1 (cf. Fig. 1) on the fly, by restricting the object class, LASER-PRINTER, to a certain price, PRICE-800,[3] which is a subconcept of PRICE (COMP-CLASS-1 = LASER-PRINTER $\sqcap$ $\forall$HAS-PRICE.PRICE-800). As a necessary result, the printer X11 is classified as belonging not only to LASER-PRINTER but also to COMP-CLASS-1. If this were not the case, either the comparison class definition or the utterance itself would be invalid. In a metarelation (CLASS-NORM-OF) the comparison class is associated with the class norm for quality, CLASS-NORM-1, which is related to the quality of the printer XII by the predicate LESS-1.

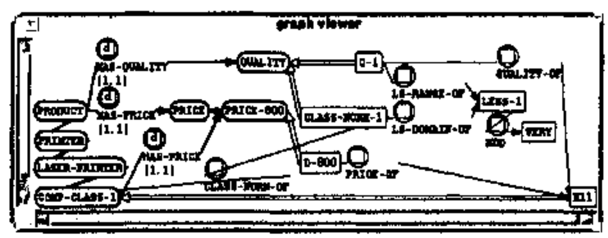

Figure 1: Representing Comparison Class and Class Norm

[2]The reasons why we focus on the interpretation of positive adjectives are twofold. First, the analytic mechanisms they require are entirely different from other kinds of degree expressions, and, second, they constitute the overwhelming proportion of degree expressions. In our corpus, e.g., we found 120 adjectives among the 4,300 words (2.8%) we considered. In [Staab and Hahn, 1997a], we analyzed comparative constructions and found approximately 150 comparatives among 30,000 words (0.5%).

[3]For reasons of simplicity of the description of the algorithm in Section 3.3, we here assume the comparison class to denote a single-valued price, rather than a more reasonable price interval.

3.2 Knowledge about Interrelations

In a discourse setting, a multitude of possibilities exist to fix the comparison class for a given adjective.

(10) Paul *is 4 years old.* He *is tall.*

(11) Paui *celebrated his 4th birthday, yesterday. He is tall.*

(12) Paui *is tall for a 4 year old boy.*

The examples (10), (11) and (12) indicate that purely linguistic restrictions are not sufficient for restricting the comparison class of an adjective. Similarly, knowledge-based computations that rely on static knowledge only, fail to determine the proper interpretations (e.g., if comparison classes cannot be created dynamically but must be predefined). Therefore, we use *(meta)knowledge of interrelations* that describes how a class subhierarchy may influence the relations of class norm instances on a scale *or* how two degrees of a given concept are interrelated. As examples, consider the sentences (13) and (14). In both of these the comparison classes are stated explicitly, and, thus, elucidate the distinction between a proper comparison class restriction and an improper one:

(13) *Peter is tall for a gymnast.*

(14) *? Peter is tall for a flute player.*

The interrelation that exists between "hasHeight" and "conductsExercises" describes gymnasts to be usually smaller than the average people. So, being tall for a gymnast does not necessarily imply being tall for the comparison class of all people. It is exactly the absence of corresponding interrelations between "hasHeight" and "playsFlute" that renders the restriction of the comparison class to flute players awkward. Two things become evident here. First, knowledge about interrelations is everyday knowledge; it constitutes a common part of human knowledge. Second, these interrelations need not be symmetrical.[4]

In order to exploit the kind of knowledge just described we use a description logics notation (as for the representation of the comparison classes mentioned before). This notation has the advantage that it combines formal explicitness with algorithmic convenience. The interrelations have in common that they describe local *restriction classes* that will later be combined to define the comparison class. We distinguish two basic types of interrelations which are illustrated by the examples (15) - (17) (the relevant comparison classes are underlined). Sentence (15) is a simple example where a degree-hierarchy interrelation is important (for a description of the relevant relations in the knowledge base (KB), cf. Fig. 2). In this example, the relevant comparison class (LASER-PRINTER) is the concept NOISE-LEVEL is *directly*

[4]Commonsense knowledge tells us that though gymnasts tend to be smaller than the average people, small people do not tend to do gymnastics very much. Assume that a population consists of 50% small and 50% tall people, respectively, 1% being gymnasts, and 90% of the gymnasts being small people. Then the probability that a gymnast is small is 90%. However, the probability that a small person is a gymnast is only 1.8%. Thus, restricting a comparison class from all people to gymnasts, in fact, decreases the class norm for height considerably, while the reverse is not true.

(15) Degree-hierarchy interrelation (with distance 1):
The noise level of the 300dpi laser printer XI1 is high for a laser printer.

(16) Degree-hierarchy interrelation (with distance 2):
The test picture of the XI1 has a good quality for the test picture of a 300dpi laser printer.

(17) Degree-degree interrelation (with distance 2):
The *XI1 offers very good quality for a laser printer that costs $800.*

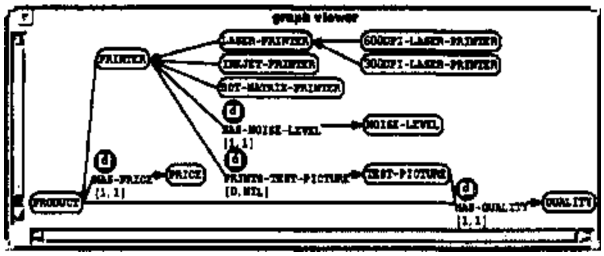

Figure 2: Hierarchy and Definitory Roles in the KB T-Box

associated with (assuming property inheritance). Therefore, the path from the relevant degree NOISE-LEVEL to the relevant restriction class LASER-PRINTER has the unit length *J* (inheritance links are not counted). (16) refers to the same type of interrelation, but differs in the length of the distance (two relations have to be passed) between one of the relevant restrictions, 300DPI-LASER-PRINTER, and the degree QUALITY (of the test picture).

In order to represent the above-mentioned interrelations knowledge must be available about which relations (in the last example: QUALITY-OF and PRINTED-BY[5]) lie between the restricting hierarchy (here, the subhierarchy of PRINTERS) and the interrelated degree (here, QUALITY). Moreover, it must be known which subclasses of PRINTER have a norm attached relating to noise level, which is either below or above the class norm associated with their direct superclass.[6] In our example, LASER-PRINTER, INKJET-PRINTER and 600DPI-LASER-PRINTER belong to the set of classes that are associated with class norms above that of their superclass, while DOT-MATRIX-PRINTER and 300DPI-LASER-PRINTER relate to corresponding lower class norms.

For degree-hierarchy interrelations we define the operator SH as in Table 2 to represent this knowledge. This operator SH takes a list of pairs of restriction classes ($RESTRC_j$) and relations (R_j). The relations, R_j, are furthermore restricted to R'_j in order to allow the definition of more constrained interrelations. This is especially necessary if the domain of a relation is not specific enough to ensure an adequate interrelation representation (e.g., cf. Fig. 3, where QUALITY-OF is restricted to REL-I-P, which can be described as *"quality-of of a test picture").* Furthermore, the

[5]For each relation (e.g., HAS-QUALITY, PRINTS-TEST-PICTURE) and relation instance we always assume the existence of its inverse which is then referred to by an intuitively plausible, name such as QUALITY-OFor PRINTED-BY.

[6]We here abstract from the consideration of multihierarchies.

$\Im_H([(\text{RESTRC}_j, \text{R}_j)], \{\text{POSC}_j\}, \{\text{NEGC}_j\})$
$:= \{\, i_1 : \text{H-INTERREL},$
$\text{R}'_1 \doteq \text{R}_1 \sqcap (\text{TOP} \times \text{RESTRC}_1), i_1\ \text{REL}_1\ \text{R}'_1, \ldots,$
$\text{R}'_k \doteq \text{R}_k \sqcap (\text{TOP} \times \text{RESTRC}_k), i_1\ \text{REL}_k\ \text{R}'_k,$
i_1 HAS-POS-CLASS $\text{POSC}_1, \ldots, i_1$ HAS-POS-CLASS $\text{POSC}_n,$
i_1 HAS-NEG-CLASS $\text{NEGC}_1, \ldots, i_1$ HAS-NEG-CLASS $\text{NEGC}_m \}$

Table 2: Representing Degree-Hierarchy Interrelations

operator takes sets of concepts that are associated with class norms above and below the class norm of their direct superclass, respectively ($\{\text{PosC}_j$ *vs.* $\{\text{NEGC}_J,\}$). SH maps them onto assertions of a degree-hierarchy interrelation instance (i_1 : H-INTERREL). These assertions are propositions *about* the relations in the T-Box, and, thus, they form an independent level of assertional metaknowledge. The knowledge required for the processing of example (16) can be recorded as in the terminological structures from Fig. 3.

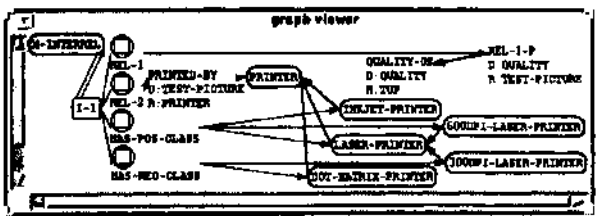

Figure 3: The Interrelation Knowledge Needed for (16).

For degree-degree interrelations (such as needed for (17)) a similar operator, SD, has to be defined. Its definition requires simpler constructs. Just like SH it takes restriction classes and relations, but instead of positive or negative subclass relations (HAS-P0S-CLASS, HAS-NEG-CLASS) its specification only includes whether an interrelation is a positive or negative correlation.

3.3 Computing Comparison Classes

Most often positive gradable adjectives refer to comparison classes that are only implicitly available (cf. (10) and (11)). In this subsection we will give an algorithm that computes implicit comparison classes by making use of semantic relations, of the knowledge of interrelations described in Section 3.2, of text-specific and world knowledge, and of the mechanism for computing explicitly given comparison classes from Section 3.1. We here assume the completion of a full semantic interpretation, including anaphora resolution.

The basic idea of the algorithm for computing comparison classes is expressed in Fig. 4: A positive adjective *a* denotes a degree *d* in the current text fragment (at present, this includes the current and the previous utterance). This degree *d* is related to an object o_1 which itself is related to another object o_p. Of course, there might be no object or several objects related to o_1, and o_p itself might have other relations as well. Each object o_i has a most specific type $C_{i,1}$. The goal of the algorithm is to select all objects o_i that are relevant for the computation of the correct comparison class. Furthermore, for each object o_i it must select its correct intermediate superconcept $C_{i,k(i)}$, which does neither restrict the comparison class too narrowly (as $C_{i,1}$ might do) nor too widely (as $C_{i,n(i)}$ might do). This goal is achieved by matching the available knowledge on interrelations against the semantic structures of the current text fragment. Finally, a comparison class is (recursively) computed by combining all the gathered restrictions. In Fig. 4, this means that the new comparison class is defined by restricting RC_1 to a new class where the role RR_p is restricted to the range RC_p.

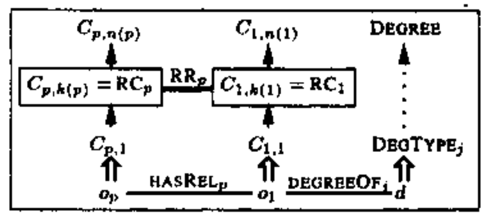

Figure 4: Relevant Structures for Computing Comparison Classes

The algorithm starts with a set of relations D_n representing the semantic interpretation of the current text fragment in which the relevant adjective a occurs.

1. Input:
   (a) A positive gradable adjective a denotes the degree d.
   (b) A set of unary and binary relations representing the meaning of the current text fragment in a Neo-Davidsonian style:
   $D_n = \{\text{DEGTYPE}_i(d), C_{1,1}(o_1), \text{DEGREEOF}_j(d, o_1), C_{p,1}(o_p), \text{HASREL}_p(o_1, o_p), \ldots\}$

Then the comparison class, COMPCLASS, is computed by applying the recursive function *"ComputeRestriction"*.

2. $(\text{COMPCLASS}, RR) =$
   $\text{ComputeRestriction}([o_1], [\text{DEGREEOF}_j])$

The main function *"ComputeRestriction"* takes a list of objects $ObjL$ and a list of relations $RelL$. The objects and relations describe the path to the degree under consideration. Thus, in the first call to *"ComputeRestriction"* $ObjL$ will always be $[o_1]$, and $RelL$ will be $[\text{DEGREEOF}_j]$.

1. Parameters given to the function:
   $ObjL = [o_m, \ldots, o_1]$, $RelL = [r_m, \ldots, r_1]$

First, *"ComputeRestriction"* tries to match[7] a piece of interrelation knowledge against the semantic structures represented by $ObjL$ and $RelL$ (step 2). If the interrelation I that has been found is a degree-degree interrelation then a singleton concept is returned (step 3). For instance, in (18),

(18) *The laser printer X11 costs $800. It offers very good quality.*

when the recursion finds the degree-degree interrelation between PRICE and QUALITY this recursion branch will return $RC = \{\$800\}$ as a singleton concept.

[7]Degree-degree interrelations have the type D-INTERREL. INSTOF, ISA, RANGE and HASROLE denote the common terminological relations. As mentioned before the relevant part of our KB is a simple hierarchy, thus there is only a single path from a concept to the topmost concept and "min" and "max" applied to a (possibly empty) subset of concepts of such a path are, therefore, partial and single-valued.

2. $(I, RR) := \iota(i_k, R_{m,k}) : |RelL| = m \wedge$ $\text{INSTOF}(i_k, \text{INTERREL}) \wedge \text{REL}_m(i_k, R_{m,k}) \wedge$ $\neg\exists R_{m+1,k}[\text{REL}_{m+1}(i_k, R_{m+1,k})] \wedge$ $\forall i \in [1, m] : \text{INSTOF}(r_i, R_{i,k})$
3. IF $\text{INSTOF}(I, \text{D-INTERREL})$ THEN RETURN $(\{o_m\}, R_{m,k})$

If a degree-hierarchy interrelation has been found then it is used to select the most specific class that fits o_m — that class is either mentioned in the degree-hierarchy interrelation or it is the maximal concept which is also defined by an inverse superrelation, R_m^{-1}, of the current relation, r_m (step 4).

4. $RC := \min_{C_k}\{C_k | \text{INSTOF}(o_m, C_k) \wedge$ $C_k \in \{C_l | \text{HASPOSC}(i_k, C_l) \vee \text{HASNEGC}(i_k, C_l)\} \cup$ $\{\max_{C_l}\{C_l | \text{INSTOF}(o_m, C_l) \wedge \text{HASROLE}(C_l, R_m^{-1}) \wedge$ $\text{INSTOF}(r_m, R_m)\}\}\}$
5. IF $RR = \perp$ THEN $RR := \iota R_m : \text{INSTOF}(r_m, R_m) \wedge$ $\text{HASROLE}(RC, R_m^{-1})$

For instance, in the first recursion for (18), LASER-PRINTER is the most specific concept of X11 that is interrelated with QUALITY and is therefore chosen as RC (the relevant interrelation is not defined here). If no interrelation were known, then PRODUCT would have been chosen, since this is the maximal concept where HAS-QUALITY is defined as a role (cf. Fig. 2 for KB definitions).

If no interrelation could be found then RC and RR must be defined in a way that renders them neutral up to the point where they themselves might become further restricted (steps 4 and 5). This is, e.g., necessary, if there are no known interrelations with relation length 1 but ones with length 2. In this case, RC must be defined "neutrally" first, and is further restricted by $RC \sqcap RR'.RC'$ afterwards. Finally, *"ComputeRestriction"* is applied recursively (step 6a), and each new restriction RC' narrows down the current restriction class RC with the terminological operation (step 6b). For instance, in the first recursion for (18) we have $RC =$ LASER-PRINTER, $RR' =$ HAS-PRICE, and $RC' = \{\$800\}$.

6. $\forall p : \text{HASREL}_p(o_m, o_p) \wedge o_p \notin ObjL$ DO
   (a) $(RC', RR') :=$ ComputeRestriction($[o_p.ObjL]$, $[\text{HASREL}_p.RelL]$);
   (where "[.]" denotes a list constructing operation)
   (b) $RC := RC \sqcap RR'.RC'$
7. RETURN (RC, RR)

The recursion stops when it finds a degree-degree interrelation or when the given semantic structures have been searched exhaustively.

3.4 Empirical Data

In an empirical evaluation we compared the algorithm presented in the previous section (henceforth, *c3*) against two simpler, naive approaches. The first of these, *n1*, uses the most specific concept of the object conjoined with the respective degree. The second one, *n2*, does not select this most specific concept, but its immediate superconcept. Interestingly, both approaches constitute somewhat of a lower bottom line for our approach, since it can switch back to one of these simpler approaches if it is unable to identify selective restrictions.

We chose a text which contained 226 sentences with about 4,300 words. 121 positive gradable adjectives were screened for which a reasonable semantic representation could be determined in 72 cases – and only these were evaluated. The remaining 59 occurrences graded idiomatic expressions, concepts that are hard to model (e.g., *"a good idea"*), or entailed other problems that were not directly related to finding the correct comparison classes. Under the assumption of complete knowledge, *c3* achieved a high success rate (60 cases (83%) were correctly analyzed). *n1* and *n2* performed much worse, as they were only able to properly determine 20 and 15 valid comparison classes (28% *vs.* 21%), respectively.

n1 and *n2* are equivalent to the procedures [Bierwisch, 1989] proposes for adjectives related with generic and non-generic nouns, respectively (e.g., in *"towers are high"* the related noun "towers" is generic, while in *"this tower is high"* it is not). An oracle that tells whether an adjective is related to a generic object and, depending on the result, changes the strategy from *n1* to *n2* would render a mechanism close to the one Bierwisch proposes. However, it would not add much benefit. Since none of the 72 considered adjectives are related to generic nouns, the positive cases of *n2* are not due to any generic use. Our results are interesting still, even though we restricted our approach to *distance-1* and *distance-2* interrelations to keep the modelling problem manageable.

3.5 Relations between Comparison Classes

Once the text understanding system has read a text and learned about properties of several objects, it will often be desirable to relate two degrees. Considerably often, these two degrees will not be related to each other but only to class norms of different comparison classes. Here, the interaction between the Degree Calculus and the determination of comparison classes comes into play.

Consider, e.g., a laser printer U13 which has been assigned a good quality comparing it to 600dpi laser printers and a laser printer V14 which has been assigned a bad quality referring to the more general class of laser printers (cf. (19)). The distance-1 interrelation describes that the quality class norm for 600dpi laser printers is higher than the one for laser printers. This statement is equivalent to an assertion like (20).

(19) $Q_{U13} \succ_0 N_{\textit{qual-600dpi-lp}} \wedge N_{\textit{qual-lp}} \succ_0 Q_{V14}$

(20) $N_{\textit{qual-600dpi-lp}} \succ_0 N_{\textit{qual-lp}}$

These pieces of information are combined by the Degree Calculus to yield the desired result, *viz.* (21).

(21) $Q_{U13} \succ_0 Q_{V14}$

A similar scheme can be given for composed comparison classes like *"laser printers that cost \$800"* and degree-degree interdependencies. Thus, knowledge about correlations does not only help to determine comparison classes; together with the representation and reasoning schemata of the Degree Calculus it even allows to bridge different comparison classes.

4 Related Work

Though comparison classes are referred to by many authors (e.g. [Klein, 1980; 1991; Simmons, 1993]), only [Bierwisch, 1989] sketches a procedure for the determination of comparison classes for gradable adjectives. However, his model is too weak to infer complex comparison classes like *"picture printed by laser printer"*. [Kyburg and Morreau, 1997] present an "extension stretching" approach that may treat referentially used adjectives, but fails to account for attributive ones which are in the focus of this paper and occur at a much higher frequency (cf. [Klein, 1979] for the referential/attributive distinction).

Inference models for degree representations usually distinguish between two major approaches. In qualitative models, such as those proposed by [Schwartz, 1989] and [Kamei and Muraki, 1994], inferences are drawn on fixed sets of two to seven ordered qualitative labels that represent degree expressions. However, while their representation schemata successfully handle several particular problems, they do not allow for the representation of comparatives and inferences like (7) or even simpler transitive inferences.

Quantitative models, on the other hand, face the principal problem that they must assign functions to lexical entries which map comparison classes to numbers. Since this assignment is highly controversial, it is usually not even done for simple examples. As mentioned before, a quantitative interval approach as proposed by [Simmons, 1993] for the linguistic description given by [Bierwisch, 1989] is unable to draw transitive inferences. The often cited advantage of interval approaches, their ability to reason with factor phrases when the input is inexact, does neither seem to be cognitively very plausible nor very important for text understanding systems. Fuzzy logic [Zadeh, 1978] faces similar problems and has been rejected with convincing arguments concerning natural language processing by [Pinkal, 1995]. However, the alternative Pinkal proposes, a supervaluation approach, though it is extremely handsome for theoretical reasons, can hardly be applied to real-world text understanding systems due to complexity considerations.

5 Conclusion

In this paper, we have defined the *Degree Calculus* the basic constructs of which (class norms, modifier labels, inference rules) are general enough to account for natural language degree expressions of a wide variety, *viz.* positives, comparatives, equatives and superlatives (in their all-quantified readings as comparatives). Furthermore, we have formulated an interpretation methodology for natural language degree expressions. Our proposal covers declarative issues of representing comparison classes and their associated class norms, as well as procedural aspects of how to determine the appropriate one for a particular adjective. The homogeneous formal treatment we give within a terminological specification framework is unique as is the provision of a comprehensive algorithm for computing comparison classes at all.

The approach we propose is based on a *deep knowledge* model of natural language understanding. We consider this a feasibility study which explores the possible limits of the terminological knowledge representation approach for an extremely demanding natural language understanding task. The empirical results we have gathered convey some good news. When we supply sufficient knowledge the approach works astonishingly well (as it outperforms naive approaches in a significant way). When we do not supply sufficient knowledge it degrades, but still produces reasonable (though suboptimal) results. We stipulate that non-knowledge-based approaches are entirely inadequate to cope with this problem at all.

Though our treatment of comparison classes is far from being complete, it is the first comprehensive treatment of comparison classes we know of. Coupled with the Degree Calculus it offers far more benefits than competing approaches.

Acknowledgements. S. Staab is a member of the Graduate Program *"Human and Machine Intelligence"* at Freiburg University, which is funded by *DFG.* We thank our colleagues, Martin Romacker and Susanne Schacht in particular, for their comments.

References

[Bierwisch, 1989] M. Bierwisch. The semantics of gradation. In M. Bierwisch and E. Lang, editors, *Dimensional Adjectives,* pages 71-261. Springer, Berlin, 1989.

[Kamei and Muraki, 1994] S. Kamei and K. Muraki. A discrete model of degree concept in natural language. In *Proc. of COLING-94,* volume 2, pages 775-781, 1994.

[Klein, 1979] E. Klein. On formalizing the referential/attributive distinction. *Journal of Philosophical Logic,* 8:333-337, 1979.

[Klein, 1980] E. Klein. A semantics for positive and comparative adjectives. *Linguistics & Philosophy,* 4(1): 1-45, 1980.

[Klein, 1991] E. Klein. Comparatives. In A. von Stechow and D. Wunderlich, editors, *Semantics,* pages 674-691. W. de Gruyter, Berlin, 1991.

[Kyburg and Morreau, 1997] A. Kyburg and M. Morreau. Vague utterances and context change. In *Proc. of the 2nd Intern. Workshop on Computational Semantics,* pages 135-155. Tilburg University, The Netherlands, 1997.

[Pinkal, 1995] M. Pinkal. *Logic and Lexicon.* Kluwer, Dordrecht, 1995.

[Schwartz, 1989] D. G. Schwartz. Outline of a naive semantics for reasoning with qualitative linguistic information. In *Proc. of the IJCAI-89,* pages 1068-1073, 1989.

[Simmons, 1993] G. Simmons. A tradeoff between compositionality and complexity in the semantics of dimensional adjectives. In *Proc. of the EACL-93,* pages 348-357, 1993.

[Staab and Hahn, 1997a] S. Staab and U. Hahn. Comparatives in context. In Proc. *of the AAAI-97,* 1997.

[Staab and Hahn, 1997b] S. Staab and U. Hahn. Dimensional reasoning with qualitative and quantitative distances. In *Proc. of the AAAI-97 Workshop on Spatial and Temporal Reasoning,* 1997.

[Woods and Schmolze, 1992] W. A. Woods and J. G. Schmolze. The KL-ONE family. *Computers & Mathematics with Applications,* 23(2-5): 133-177, 1992.

[Zadeh, 1978] L. A. Zadeh. PRUF — a meaning representation language for natural languages. *International Journal on Man-Machine Studies,* 10:395-460, 1978.