Adjectival vs. Nominal categorization processes: The Rule vs. Similarity hypothesis*

Galit W. Sassoon, ILLC, University of Amsterdam

Abstract

Classification of entities into categories can be determined based on a rule – a single criterion or relatively few criteria combined with logical operations like 'and' or 'or'. Alternatively, classification can be based on similarity to prototypical examples, i.e. an overall degree of match to prototypical values on multiple dimensions. Two cognitive systems are reported in the literature to underlie processing by rules vs. similarity. This paper presents a novel thesis whereby adjectives and nouns trigger processing by the rule vs. similarity systems, respectively.

The paper defends the thesis that nouns are conceptually gradable and multidimensional, but, unlike adjectives, their dimensions are integrated through similarity operations, like weighted sums, to yield an overall degree of match to ideal values on multiple dimensions. By contrast, adjectives are associated with single dimensions, or several dimensions bound by logical operations, such as 'and' and 'or'. In accordance, nouns are predicted to differ from adjectives semantically, developmentally, and processing-wise. Similarity-based dimension integration is implicit – processing is automatic, fast, and beyond speaker awareness, whereas logical, rule-based dimension integration is explicit, and is acquired late.

The paper highlights a number of links between findings reported in the literature about rule- vs. similarity, and corresponding structural, distributional, neural and developmental findings reported about adjectives and nouns. These links suggest that the rule vs. similarity (RS) hypothesis for the adjective-noun distinction should be studied more directly in the future. The implications of the hypothesis are examined against paradigmatic and non-paradigmatic nouns and adjectives, including nominalizations, animate-evaluative nouns, relational adjectives and determiner-adjective constructions.

1 The noun-adjective puzzle

1.1 Preliminaries: Dimension integration in nouns and adjectives

Scholars generally agree that natural languages provide evidence for a typology of predicates consisting of word classes such as nouns, adjectives and verbs (Baker 2003: 1-16).[1] Formal semanticists, using classical logic or similar logical systems, typically classify these word classes uniformly as 'predicates', disregarding typologically dominant morpho-syntactic, and semantic-pragmatic differences. Some semantic analyses distinguish between verbs and other predicate types, analyzing verbs as denoting event types, but they do not say what distinguishes between adjectives and nouns. Nouns tend to occur in argument position, where their main function is to refer to objects, whereas adjectives typically occur in predicate- or modifier-position. Nonetheless, unlike verbs, adjectives rarely inflect for tense and aspect, and in many languages for adjectives to occur in predicate position they have to combine with a copula or an affix. Hence, on the Kamp-Montague approach, adjectives are not predicates, but predicate modifiers (Kamp 1975). One problem with this approach as a basis for the noun-adjective distinction is that nouns can freely modify nouns too, as in *elephant pig* and *brick factory* (Baker 2003, Chap. 4).

Intuitively, people suppose that nouns, like *bird*, denote 'object categories', while adjectives, like *red*, denote 'properties'. Psycholinguists employ this distinction, but do not explicate to what exactly it reduces. It certainly does not parallel the formal

distinction between extension and intension, which for predicates reduces to an entity set vs. a property, respectively; after all, *red* can refer to the set of red objects – the extension of *red*, and *bird* – to the property of being a bird – the intension of *bird*. This paper aims to explicate what the intuitive difference between properties and categories amounts to.

Grammatical categories have a neuro-anatomical basis; for example, Miozzo et al. (2010) review reports about bilingual aphasics involving disruption selectively affecting the production of verbs vs. nouns, and regular vs. irregular verbs. Critically, these selective deficits were manifested in a strikingly similar manner across the two languages spoken by each of the individuals. Cappelletti et al (2008) show that the left prefrontal cortex is selectively engaged in processing verbs. These authors write: "*Other word classes, like adjectives, may also be associated with distinct neural circuits, but the status of categories other than nouns and verbs has not been well studied."* (Cappelletti et al 2008, note 1, P. 718). This situation originates in part due to the fact that nouns and verbs are more distinguished in human languages than adjectives are (Evans 2000; Sapir 1921).

The various aspects of lexical knowledge, including semantic characteristics ('dimensions', such as feline, pet and furry for *cat*), syntactic features (noun, countable) and morphological ones (type of plural inflection, etc.) are represented and stored with sufficient independence (Miozzo et al. 2010) – brain injury may disrupt one of these aspects while sparing others (Rapp and Goldrick 2006; Hillis and Rapp 2001; Miceli et al. 2002). Importantly, syntactic categories do not align with semantic dimension types (for example, abstract vs. visual vs. motor dimensions). Kellenbach et al (2002) used event-related potentials to investigate three subclasses of nouns and verbs, which differed in the type of salient semantic dimensions. ERP effects were observed for both grammatical class and dimension type, with no interactions between the two. The effects of grammatical class (verbs vs. nouns) did not differ significantly between the semantic dimension types, and, conversely, the effects of semantic dimension type (abstract, visual, and motor) were equivalent for each grammatical class. Thus, lexical-semantic knowledge is organized in a manner that takes account of grammatical category and dimension-type distinctions, independently (Kellenbach et al. 2002: 564-5).

What, then, do word class distinctions represent? The present study investigates the proposal that, rather than on the type of dimensions in their dimension list, the difference between nouns and adjectives hinges on the default way the dimensions are integrated. The basis for this idea comes from the cognitive-psychological study of categorization and learning of artificially construed categories. On a first sight, the huge literature on these topics often appears to report inconsistent results. However, a consistent picture is revealed when studies are divided by categorization type (Ashby and Maddox 2005), namely by the way separate semantic dimensions are integrated to form categorization judgments.

Theoretically, the set of instances denoted by an adjective or a noun is independent of any specification of dimensions (Kripke 1972; Putnam 1975). However, empirically, specification of dimensions has been found to be persuasive and helpful for speakers, who have to constantly resolve categorization tasks within contexts of partial information about the extensions of predicates. In addition, the ways dimensions are structured to yield categorization criteria affect language processing and use. They have semantic implications, for example, on inferences (Murphy 2002: chap. 8 and 11), as well as morpho-syntactic implications. This paper introduces the thesis that dimensions have to be considered by an analysis of the adjective-noun distinction. In other words,

the ways by which speakers use dimensions to determine what the reference of a predicate is, is sensitive to the adjective-noun distinction.

This paper treats nouns and adjectives in language, and their correlates in thought, on a par, on the basis that underlying conceptual structures – dimension sets – are needed to account for the linguistic phenomena under discussion (for example, the licensing of dimension arguments and degree morphemes). A distinction between language and thought can be promoted by separating between dimensions which are specified in the lexicon and ones which are merely part of an encyclopedic memory of general world knowledge separate from the semantics of words. This paper is consistent with such an approach, but is not involved in such a project, which requires proper justification as to where each dimension belongs.[2]

Two important types of categorization, which have different neural and developmental correlates, include categorization based on rules vs. similarity (Ashby and Maddox 2005). Classification in rule-based categories depends on a single dimension or a simple enough conjunction or disjunction of dimensions – one that people can reason about explicitly. In accordance, processing is by an explicit, declarative memory system, namely, a system for storing and retrieving memories we are aware of and can declare about. Conversely, in similarity-based categories, information about degrees of instances in multiple dimensions is integrated by averaging, or in a holistic 'gestalt' manner. These types of information integration occur at an early processing stage. In accordance, the dimensions and the way they integrate to create categorization judgments are hardly accessible through introspection. Processing is accomplished through an implicit, procedural memory system, namely it is reflexive and automatic, like, for example, the capacity to drive.

This paper explores the hypothesis that the morpho-syntactic cues distinguishing between adjectives and nouns trigger processing by the rule vs. similarity-based categorization systems, respectively. Logically speaking, noun categories can be structured by means of rules, and adjectival categories can be structured by means of similarity. The thesis of the paper is that these structures are possible, but are not the default, prominent ways to resolve reference of nouns and adjectives. The proposal is tested against linguistic data, as well as experimental data from psycholinguistic studies that indirectly bear on the issue. Factors that can interact with a word's syntactic category to yield non-default processing are discussed, as well.

Finally, this is not meant to be a study whereby a concept is analyzed in terms of simpler primitives; rather, dimensions of predicates are viewed as predicates on their own right, either simpler or more complex than the predicates they are dimensions of. The crux is that the interpretations of predicates and their dimensions are mutually constraint in a systematic way, which is determined by the nature of the relation 'dimension of'. Therefore, to set up the ground for the formalization of the hypothesis, we have to specify the ingredients of interpretation of predicates, and the relation 'dimension of' between predicate interpretations.

Importantly, there are many possible 'dimension of' relations. The distinction between categorization by rule and by similarity amounts to categorization along two different types of 'dimension of' relations – two principally different ways in which dimensions can be integrated to create a unified categorization criterion. Similarity based categorization criteria drawn from psychological theories of concepts are introduced below, and their extensive empirical support and relevance as a basis for an analysis of nouns is reviewed in section 1.3. Rule-base categorization is also introduced shortly, and is supported as a basis for an analysis of adjectives in section 1.4, and later on in sections 2 and 3.

1.2 The Rule vs. Similarity (RS) Hypothesis for the Adjective-Noun Distinction

The main idea defended in this paper is that the noun-adjective word class distinction probes similarity vs. rule based categorization – categorization along different types of dimension-of' relations.

A rule-based dimension-of relation amounts to integration of dimensions through logical operations, such as those denoted by *and* and *or*. For example, the dimensions $P_1 \ldots P_n$ of a rule based category P can be integrated through dimension conjunction. In this case, $P_1 \ldots P_n$ are dimensions of P iff for any individual x, x falls under P iff x falls under ALL of P's dimensions: $P(x) \Leftrightarrow P_1(x)$ & …. & $P_n(x)$. Alternatively, dimensions can be integrated through dimension disjunction, in which case $P_1 \ldots P_n$ are dimensions of P iff for any individual x, x falls under P iff x falls under SOME of P's dimensions: $P(x) \Leftrightarrow P_1(x)$ or …. or $P_n(x)$. Furthermore, x falls under a gradable rule-based predicate P iff the degree to which x exemplifies P, deg(x,P), exceeds a membership threshold, s(P) (henceforth, P's standard), and gradable predicates $P_1 \ldots P_n$ are dimensions of a conjunctive rule-based predicate P iff for any individual x, x's degree exceeds the standard in ALL these dimensions: $P(x) \Leftrightarrow deg(x,P_1) > s(P_1)$ & …. & $deg(x,P_n) > s(P_n)$.

For example, everybody will tell you, as a simple matter of language use, that if you are healthy, you cannot have any serious disease, whereas if you are sick, you must have some disease or other. Therefore, *healthy* means healthy in all dimensions, and *sick* – sick in some dimension. We may consider one to be healthy despite, say, high blood-pressure only when this dimension is considered irrelevant. When using expressions like *all* or *everybody*, the standard practice is to ignore irrelevant entities (von Fintel 1994), but not to allow any other exceptions.

Hence, I propose that adjectival dimensions integrate through logical operations like those denoted by *all* and *some* (cf. Stenning et al. 2008). Thus, x is healthy holds true in a context c whereby F(healthy,c) is the dimension set of *healthy* iff $\forall F \in F(healthy,c)$, $deg(x,F,c) > s(F,c)$; in words, x is healthy with respect to every dimension F in F(healthy,c), meaning that x's degree exceeds the standard of membership of every dimension in c, e.g., blood pressure AND cholesterol AND sugar, etc. By contrast, x is sick holds true iff $\neg\forall Q \in F(healthy,c)\ deg(x,F,c) > s(F,c)$; in words, x is sick with respect to at least one dimension, meaning that x fails to exceed the membership standard of blood pressure OR cholesterol OR sugar, etc.

Turning to similarity-based classification, the most influential analysis of linguistic concepts in the psychological study of concepts is called the prototype theory. The origins of the prototype theory go back to Wittgenstein (1968 [1953]) in "Philosophical investigations", and within cognitive psychology it has become a unified approach due to the extensive work of figures like Eleanor Rosch, Tversky and their associates.[3] The most central alternative to the prototype theory is the exemplar theory. [4] By and large, other cognitive approaches can be seen as branches of these two theories, including the similarity-as-transformation approach, which for linguistic concepts converges with standard similarity-based theories (Chater and Hahn 1997: 33), and the knowledge approach (Murphy 2002: Chap. 6), which makes use of dimensions based on elaborate encyclopedic knowledge. For reasons of space and clarity, this paper focuses mainly on prototype accounts, but the conclusions it draws are consistent with the other approaches.

A similarity based dimension-of relation amounts to integration of dimensions through non logical, similarity operations, such as weighted sums and products. For

example, the dimensions of a similarity-based category P can be integrated through a weighted-sum function. In this case, $P_1 \ldots P_n$ are dimensions of P and $W_1 \ldots W_n$ are their respective weights, representing their importance for categorization under P, iff for any individual x, x falls under P iff the weighted sum of x's degrees in P's dimensions exceeds P's standard: $P(x) \Leftrightarrow W_1 deg(x,P_1) + \ldots + W_n deg(x,P_n) > s(P)$. An important difference between this similarity-based categorization criterion and a conjunctive rule-based categorization criterion is that, in the latter, the degrees of an instance have to exceed the standard in each and every dimension, whereas in the former, entities may fail to fall under some dimensions or others. It suffices that their degrees in the dimensions add up to a high enough overall degree.

In the above mentioned psychological studies of similarity-based categorization, dimensions are often selected based on lists provided by participants, and dimensions are weighed in various ways, for example, depending on the frequency at which speakers list them as dimensions. However, it has also been shown, through simulations with connectionist networks, that dimensions and dimensional weights can be gradually learnt through feedbacks on correct and incorrect categorization judgments. Generally, the extension can be learnt from the dimensions, and vice versa. These studies illustrate that the contextual extension of a similarity-based predicate is constrained to instances with high enough weighted sum of degrees in the dimensions. Significantly, no criterion for dimension selection is needed other than this categorization constraint. Children can use it to extract dimensions from evidence concerning category instances, and latter on they can use the extracted dimensions to classify newly encountered instances.

Moreover, similarity-based categorization is rarely represented by the mere degrees of entities in the dimensions, $deg(x,P_i)$; rather, normally, it is represented by the distance between these degrees and respective ideal values for the given category in the given dimensions, $deg(P,P_1) \ldots deg(P,P_n)$. These values represent the prototype of P, but no actual entity should instantiate all of them. The distance of instances x from the prototype of P, D(x,P), equals the weighted sum of their distances from the ideal values on the dimensions: $D(x,P) = W_1|deg(x,P_1) - deg(P,P_1)| + \ldots + W_n|deg(x,P_n) - deg(P,P_n)|$. This distance should be overall small enough. Hence, for entities to count as P, their overall similarity to the prototype of P, deg(x,P), has to exceed P's standard, S(P), where an entity's similarity is inversely related to its overall distance from P's prototype, D(x,P). According to Shepard's (1987) *Universal Law of Generalization*, similarity is an exponentially decaying function of distance: $deg(x,P) = 2^{-D(x,P)}$. This law applies across a wide range of stimuli, and applies both to people and to animals. This renders a similarity-based categorization judgment P(x) equivalent to $2^{-D(x,P)} > s(P)$ (Murphy 2002; Chater and Hahn 1997), which reduces to $deg(x,P) > s(P)$.

Let us illustrate with a concrete example. Consider, for example, a psychological analysis of nouns like *bird;* experiments suggest that the conceptual structure of this noun includes a rich cluster, F(bird,c), of dimensions, which is called *the prototype* or *summary representation* of the concept. It include features like 'bird-genotype', 'bird descendant', 'can interbreed with birds', or merely 'winged', 'feathered' and 'small'. Each dimension F in F(bird,c) has an attention weight W_F. For example, W_{size} tells us how important size is in discriminating birds from non-birds. In addition, the conceptual structure of *bird* includes an ideal value on each dimension; for example, deg(bird,size,c), represents the ideal size for birds.

In addition, a similarity structure includes a mapping of entities x to degrees, deg(x,bird,c), representing their similarity to birds, the extent to which their values on

the dimensions match the ideal values for birds or bird-types (Murphy 2002: chap. 3; Hampton 1995, 1998).

Moreover, the degrees of an entity in the different bird dimensions integrate into a unique degree in the given noun by means of similarity functions. Mathematically speaking, most of these are averaging operations (Sassoon 2011a), e.g. weighted-sums, as in the example above, or weighted-products (Hampton 1995; Murphy 2002: 68). Degrees are modeled on a scale between 0 and 1; dimension weights are all positive and sum up to 1. Averaging on dimensional degrees captures the fact that, by and large, no single dimension is a necessary or sufficient condition for falling under a noun (Wittgenstein 1953); rather, the membership condition requires that the average degree of entities in the dimensions exceed a certain threshold (for supporting discussion see section 1.3). Thus, abstracting away from fine-grained distinctions between different psychological models, x is categorized as a bird iff x's degree of birdhood exceeds the standard: deg(x,bird,c) > s(bird,c); x's degree of birdhood represents the degree of similarity of x to a bird, by virtue of being inversely related to x's average distance from the ideal bird values on the dimensions in F(bird,c): $D(x,bird,c) = \Sigma_{F_i \in F(bird,c)} W_i |deg(x,F_i,c) - deg(bird,F_i,c)|$.

In certain cases, categorization proceeds through categorization under a sub-category of the category; e.g., an entity x can be classified as a bird either because it sufficiently resembles the prototype – ideal values on the dimensions – of a bird – deg(x,bird) > s(bird) – or because it sufficiently resembles the prototype of one of the bird exemplars – a robin, an ostrich, a chicken, etc. Categorization may even be based on similarity to encoded instances of the category. The same basic mechanism represents all these cases. Exemplars are categories on their own right, which are associated with dimensions, degree functions and so on. It is, therefore, not very surprising that exemplar-based categorization is prevalent. Instances provide sets of dimensions and ideal values which are their own values on salient dimensions under which they fall.

In the case of two or more prototypes, as in exemplar-based categorization, entities are classified under the category to which prototype they resemble most: P(x) iff for all contrasting categories Q, deg(x,P) > deg(x,Q).[5] For example, considering the contrast categories *birds*, *mammals* and *repltiles*, x is classified as a bird iff x's similarity to the prototype of bird, deg(x,bird,c), is bigger than x's similarity to the prototype of *mammal* or *reptile*.

Notice that one-dimensional, rule-based categories may also relate to degrees representing distance from ideal values; for example, intuitively, an entity x counts as healthy with respect to blood pressure iff x's blood pressure is close enough to an ideal value, i.e., x's degree falls within a normative range surrounding an ideal. Hence, the fact that in similarity-based categories gradability is based on distances from ideal values does not distinguish them from rule-based categories; rather, the fundamental difference between similarity-based and rule-based categories lies in the way the dimensions are integrated. In the former, dimension integration is based on similarity functions, which typically involve averaging, and in the latter it is rather based on logical operations, such as ALL and SOME.

In sum, the name *similarity* highlights reference to an ideal, but the defining characteristic binding the multiple and variable accounts within the similarity approach is the rejection of logical dimension integration, in favor of non logical functions such as weighted sums and products (averaging). In fact, the similarity based approach came into being as a response to 'the classical theory', namely a family of accounts that are bound together by the use of logical dimension-integration

operations, such as conjunctions and disjunctions (Lakkof 1987: Chap. 1; Hampton 1995; Murphy 2002: Chap. 1).

We have by now specified all the ingredients of predicate interpretation. Formally, in each context c, each predicate P, whether a noun or an adjective, is associated with:

(i) a degree function, $\lambda x.deg(x,P,c)$, i.e. a mapping of entities x to degrees deg(x,P,c);

(ii) a standard of membership, s(P,c), such that entities are P iff their degree in P exceeds P's cutoff point in c, and

(iii) a set of predicates, F(P,c), P's dimensions in c, and their weights, and the ideal values for P, $W(P,F_i,c)$ and $deg(P,F_i,c)$, respectively, for each dimension F_i in F(P,c).

On this model of interpretation, all predicates are gradable in the sense of association with a degree function, dimensions, and standard, pace the more standard approach, whereby gradability (or its absence thereof) constitutes a core semantic difference between adjectives and nouns (Kamp 1975; Kennedy 1999). Yet another existing approach associates nouns and adjectives with the same type of semantic structures – the differences between them are derived from mere syntactic features, which are on the whole void of semantic content (Baker 2003: chap. 1). By contrast, on the present approach adjectives and nouns are not semantically alike. The ways their dimensions are glued together differ fundamentally, resulting in rule vs. similarity categories. Sections 1.3-1.4 support the cognitive realism and linguistic relevance of gradability and dimensionality in the noun domain, as well as the proposed difference between nouns and adjectives.

For the most parts of the paper, the context indices c are omitted, but not because the ingredients of interpretation do not vary between contexts. In nouns, speakers tend to encode a set of dimensions which they use in order to categorize new entities in new contexts. When an entity set can be indicated by several different combinations of similarity dimensions, different conventions may prevail in different populations, scientists vs. laymen (cf. section 1.3), etc. Also the importance of a dimension may vary to a certain extent, depending on contextual purposes, and so does the standard of membership, and in accordance the extension. However, variance of the latter is relatively limited. Finally, the tendency to employ a prototype or exemplars appears to vary across concepts, contexts, and speakers (Smith and Minda 1998; 2000). From the semantic perspective, the tendency aligns with interpretation. The interpretation of a noun like *bird* may consist of realizations of the kind 'bird' or of bird sub-kinds (Carlson 1977; Dayal 2004; Sassoon 2007: chap. 8).

Moreover, in adjectives, the dimensions are highly context dependent. Importantly, no criterion for dimension selection is needed other than the categorization criterion – x is healthy $\Leftrightarrow$ For ALL F in F(healthy,c), $deg(x,F) > s(F)$. Using such criteria, children in acquisition, or adults within new contexts, can extract category dimensions from evidence about category instances, and latter on use these dimensions to classify newly encountered entities. However, as observed by Lewis (1979) and von Fintel (1994), among many others, natural language quantifiers are by default restricted to those entities that count in each context of use. For example, intuitively, if I have a heavy cold, I am not healthy, but if the issue is whether I can embark on a glacier ski tour this coming spring, I may count as healthy. In considering the spring vacation, a passing cold is ignored. Some contexts require high standards of precision and then it is more difficult to discard dimensions as unimportant or irrelevant; e.g., normally, one can be considered *healthy* despite a slight cold, but not in a context of a pre-surgery medical examination, whereby ALL dimensions, even unimportant ones,

count. One would not check 'healthy' upon filling a medical questionnaire in such a context, because one is not strictly speaking *healthy*.

Context restrictions on quantification domains are prevalent. For example, the domain of the universal quantifier denoted by *everything* in (1a) clearly does not include every possible object, only sites in Paris, or maybe even only famous or adored sites. Consider also the example in (1b), which the secretary in the linguistics department at Tel Aviv University posted on the bulletin board. Naturally, notices posted by the secretary herself are excluded. Finally, the generic statement in (1c) and conditionals in (1d-e) illustrate the role of restrictions in 'encyclopedic' knowledge of scientific generalizations, as well as in daily life reasoning. The conditional in (1d) is restricted to eventualities in which there is no oil in the tea. This is crucial to account for the fact that, intuitively, (1d) fails to entail (1e).

1) a. I lived near the Seine, near Boulevard St. Germain and Rue St. Michel, near the market and the Pantheon, near everything.
   b. No notice (of any kind) may be posted on the notice board.
   c. A duck lays whitish eggs
   d. If there is sugar in the tea, the tea tastes well.
   e. If there is sugar and oil in the tea, the tea tastes well.

Conventionally, contextual restrictions on quantifying expressions are represented by a context variable (say-, X_c), whose value is a set of relevant individuals (von Fintel 1994). The truth conditions of a statement with a quantifying expression, like (1c), are considered as conveying that every individual, which is a duck, and is in the set of the relevant individuals, lays whitish eggs. The context variable is often viewed as a set of properties that contextually restrict the domain (Kratzer 1979, 1986; Kadmon and Landman 1993; von Fintel 1994). The context set for (1c) may include properties like *water bird, of the family Anatidae, with webbed feet and flattened beaks,* etc., which usually characterize ducks, but it may also include properties like *adult* and *female.* When considering reproduction, being a *female* is a necessary condition for relevance; if male ducks were relevant, the statement would be judged false, and it is not. Addressees normally grant that the preconditions for, e.g., laying eggs are satisfied, leaving open only the question under discussion – whether the eggs are whitish or not (Partee 1989). The predicate itself – *laying whitish eggs* – is not regarded as constraining the domain, for this would render the truth of a generic trivial. Additional systematic constraints on domain restriction, discussed in Partee (1989), depend on intonation and focus.

Hence, natural language quantification is, by default, restricted, in relatively systematic ways, and quantification over dimensions is no exception. Statements of the form *x is healthy* convey that "x is healthy in every respect that isn't ignored". The set of dimensions can be restricted by all kinds of properties of health dimensions, such as *stable* (to exclude, e.g., a passing cold), *important* (to exclude, for instance, a weak Asthma), and so on, but it cannot be restricted by, for example *properties with respect to which x is healthy*.

To summarize, I hypothesize that the adjective-noun distinction probes processing by two different underlying cognitive systems, yielding different types of dimension integration:

2) The Rule vs. Similarity (RS) hypothesis for the noun-adjective distinction:
   a. *Nouns denote similarity-based categories*: They are associated with multiple dimensions (characteristic features) which are integrated through non logical, similarity operations, such as averaging, at an early processing stage; for an entity to be classified under a noun, its average similarity, e.g., its weighted sum or product on the dimensions of the category or of one of its exemplars, should reach the membership standard.
   b. *Adjectives denote rule-based categories*: They are associated with either a single categorization criterion, or a set of criteria integrated through logical operations, like conjunction and disjunction (or the respective quantifiers); thus, to count as an instance of an adjective, an entity has to reach the standard in either a single dimension or a dimension-conjunction or -disjunction. Processing of dimensions is explicit.

The basic motivation for the RS hypothesis is that it provides an attractive solution to the category vs. property puzzle related with the noun-adjective word class distinction. Why are nouns and adjectives intuitively seen as denoting 'object categories' and 'properties', respectively? Each object corresponds to a plurality of dimensional values. According to the rule vs. similarity (RS) hypothesis, more features of denoted objects are encoded as part of the interpretation of nouns than of adjectives. Thus, in designating overall similarity of objects in multiple dimensions, nouns are more readily associated with the denoted objects than adjectives are. What is more, according to the RS hypothesis, in usage of adjectives, but not nouns, speakers are aware of the dimensions. Thus, speakers experience nouns as directly designating object sets, while adjectives are experienced as designating single properties of objects. According to this new perspective, formal semanticists are correct in unifying the semantics of nouns and adjectives, both eventually denote object sets.[4] At the same time, the intuitive category vs. property distinction suggest that intuitive differences reside in whether dimension integration involves explicit or implicit processing.

A variety of additional consequences of the RS hypothesis are supported empirically. The basic prediction of the RS hypothesis is the following.

3) a. If nominal dimensions integrate through similarity, they do not add categorization criteria.
   b. If adjectival dimensions are rules, they function as categorization criteria.
   c. The different ways degrees are determined in adjectives and nouns yield differences in the range of degree modifiers they license.
   d. Explicit processing of the dimensions of adjectives, but not nouns, triggers licensing of an overt dimension argument and quantification over dimensions.
   e. The rule and similarity structures underlying adjectives and nouns yield neural and developmental differences.

The rest of part 1 surveys evidence to the effect that predictions (3a-d) are borne out. Classification in adjectival and nominal categories is based on rules vs. similarity, respectively. This distinction explains significant semantic differences between adjectives and nouns. Prediction (3e) is discussed in part 2; it is shown that that adjectives and nouns differ developmentally and neutrally. The differences are consistent with the RS hypothesis. Part 3 shows that the RS hypothesis fairs well,

considering a number of non-paradigmatic sub-classes of nouns and adjectives in English and other languages, such as evaluative nouns (like *genius*), nominalizations (like *similarity* and *success*), and so on.

1.3 Similarity – averaging-based – classification in nouns

Experimental results established that speakers characterize nominal concepts (nouns and noun phrases) by a rich set of dimensions. For example, the noun *bird* is characterized by dimensions like feathers, small size, flying, singing, perching, eating insects, etc. (e.g., Rosch 1973).

Wittgenstein (1968 [1953]) and Fodor et al. (1980) have shown that the dimensions that speakers associate with nominal concepts are not definitional. They do not stand for necessary conditions for membership in the denotation. Their arguments were supported in experiments with many different types of nouns (Hampton 1979 and 1995). Naturally, abstract words such as, for instance, *happiness*, *war*, *identity*, *difference* and *possession* are typically characterized by complex clusters of non-definitional, characteristic features. Also, it is relatively easy to see that no clear cut criteria distinguish between chairs, armchairs and sofas, or between glasses and vases. Yet, it has also been shown that, for example, kinship categories, which can be defined precisely, are typically processed using multiple similarity dimensions, not categorization criteria; e.g., the category *grandmother* could be processed as merely equivalent to *mother of a mother*, but in practice people use a significantly richer cluster of perceptual and behavioral features, which do not stand for necessary conditions for membership. Nor do the noun dimensions stand for sufficient conditions for membership. For example, popes and homosexuals who have been living with a partner for many years are adult males that were never married, but are they *bachelors*?

Even the interpretations of names of natural-kind concepts, such as *horse*, *tiger* and *dog*, and names of substances, such as *water*, *gold* and *metal*, are affected by dimensions, which fail to be strictly necessary and/or sufficient for membership. Furthermore, the nonexistence of a definition is by no means constrained to ordinary usage; it has been shown that experts such as doctors, X-ray technicians and biologists use symptomatic rather than definitional features to characterize, e.g., disease types, x-ray findings and species, respectively (Murphy 2002: Chap. 1).

For example, a horse genotype is intuitively and scientifically thought to be necessary for horses. Yet, experiments show that entities that possess, for instance, a zebra genotype, but due to a special diet or medical treatment become highly similar to horses in appearance and behavior, are often judged by laymen to be horses. While the convention held by laymen may be considered merely wrong, experts use sophisticated, scientifically established classification methods to justify some taxonomy of species or other. However, natural kind terms still escape strict definition; for example, the relatively solid criterion of interbreeding, whereby each natural kind covers a set of instances that can interbreed, leads to a gamut of mutually overlapping species; see van deemter (2010; chap. 2) for a detailed discussion of the insufficiency of this and other scientific criteria for natural kinds terms.

Similarly, metallurgists define the substance name *metal* to be an element that has metallic properties, such as electrical conductivity, thermal conductivity, ductility, malleability, strength, and high density. However, the conceptual and semantic structures underlying the word *metal* do not guide speakers, including experts, as to whether all or some of the dimensions have to hold for an element to classify as a

metal. In accordance, metallurgists fail to get to agreement about the number of properties that have to hold – some say three, some say six, etc. They have no other option but to proceed presuming that they are talking about the same thing (Murphy 2002: 18).

However, empirical research of concepts suggests that the dimensions of a noun are intersubjective and convention based, and so is a noun categorization criterion. The noun dimensions are not bound through quantifiers such as *all*, *some* or *three*, which answer questions such as how many dimensions have to hold of instances; rather, it is through similarity structures that the noun dimensions help to indicate membership, as well as typicality (exemplariness). Concept-theories, which analyze nominal concepts through averaging (weighted sums or products), are extensively supported; for an extensive review see Murphy (2002: chap. 3) and Hampton (1998). These theories have been introduced in section 1.2. Briefly, the *distance* of an entity x from the ideal value for a category P on a dimension F, D(x,P,F), is the difference between x's and P's values on F (cf. (4a)). If x and P completely match on a dimension the distance is 0. Otherwise, the distance may in principle be infinite, but it is usually modeled on a 0 to 1 scale. If dimensions are treated as non-gradable (mapping entities to either degree 1 or degree 0), distance in a dimension is, accordingly, either 0 or 1.

4) a. The distance of x from the ideal value for P on dimension F:
$$D(x,P,F) = |\, deg(x,F) — deg(P,F) \,|.$$
b. Arithmetic mean-distance (for $F(P) = \{F_1, \ldots, F_n\}$):
$$D(x,P) = W_{F1}D(x,P,F_1) + \ldots + W_{Fn}D(x,P,F_n).$$
c. A categorization criterion: $x \in [[P]]$ iff $2^{-D(x,P)} > s(P)$.

The *average distance* of x from P (in the concept's dimension set), represented in (4b) using an arithmetic mean, is the sum of x's weighted degrees in every dimension; the dimension weights are all positive, and they sum up to 1. The degree of similarity of x to P, deg(x,P), is assumed to be inversely related to x's distance from P, D(x,P).

The similarity degrees predict judgments of exemplariness, likelihood and speed of classification, speed of retrieval, accuracy of memory, etc.

First and foremost, the last forty years of research in cognitive psychology have established beyond doubt that speakers consider certain entities as better examples than others of concepts that nouns denote. For example, robins are considered more typical or representative of birds than ostriches, and bats are considered more related or similar to birds than cows. The averaged similarity degree of an item on the concept's dimensions is a good indicator of its typicality. For example, a robin scores higher, on average, than an ostrich, on *small, flying, perching, singing, eating insects*, etc. In addition, the overall similarity of an item to contrasting concepts is inversely related to its typicality in the concept. For example, if two items have the same overall similarity degree on the *bird* dimensions, the one with the lower degree of similarity to optimal values on the dimensions of, for instance, *mammals*, is regarded as a better example of a bird (Rosch and Mervis 1975).

Hence, by identifying typicality with similarity to the prototype, the prototype approach derives the basic typicality effects, namely, the fact that speakers order entities by typicality and the fact that the noun dimensions are ordering dimensions, which together help mapping entities to degrees.

Second, Hampton (1998) has found a very strong coupling between the typicality ratings and the ratings of membership-probability, in about 500 items of 18 concepts denoted by nouns (henceforth, nominal concepts). Moreover, any deviations from this

pattern were highly systematic. Thus, in nominal concepts, typicality and (subjective) membership likelihood tend to go together. Moreover, these experimental findings suggest that nouns have borderline cases. Significantly, judgments about the concept membership of, for instance, curtains for *furniture* or avocado for *vegetables*, are much less consistent then judgments about clear instances (McCloskey and Glucksberg 1978; Hampton 1998; Murphy 2002).

Finally, typicality judgments are connected to numerous processing effects. Most importantly, typicality correlates with online categorization times. For example, when robins are considered better examples of birds than ostriches, verification time for sentences like *A robin is a bird* is faster than for sentences like *An ostrich is a bird* (Rosch 1973; Roth and Shoben 1983). Retrieval of concept instances from long term memory is performed by means of serial search, which begins with the best examples of the concept (Rosch 1973 : 140-141).

The prototype approach assumes that categorization under nominal concepts is based on similarity to the prototype. A certain similarity degree forms the categorization standard. An entity is classified as P iff its similarity degree reaches this standard. In other words, categorization is a process in which it is decided whether the average distance of an entity from the prototypical dimensional values is small enough, cf. (4c). This standard-based categorization-principle accounts for many typicality effects. First, it predicts the fact that likelihood of categorization is by and large monotonically related to similarity to the prototype.

Second, the existence of borderline cases – a gap – in nouns is predicted, because for some entities it may not be clear whether they reach the standard or not, if their degree is very close to the standard, or if the convention specifies the standard only approximately, cf. the vagueness concerning the cutoff between chairs and benches, plates and bowls, etc. Since nouns clearly exhibit properties of vague concepts (in particular, different forms of the Sorites paradox), this feature of the analysis is also a basis for an account of vagueness effects in the nominal domain. This is important as most formal theories of vagueness focus primarily on adjectives (cf. Kamp 1975).

Third, by assigning important dimensions (like "horse genotype") a particularly high weight, this principle derives the intuition that they can almost count as necessary and sufficient for membership. They might be violated only if the values on other dimensions are sufficient to reach threshold (Hampton 1979).

Fourth, the above mentioned online processing effects are predicted, too. In items with low degrees in some dimensions, more dimensions need to be considered in order to determine that they reach the threshold for membership.

Last but not least, the set of known concept members plays a crucial role when the concept standard is unknown. A similarity-based categorization criterion predicts that newly encountered entities, whose overall similarity is higher than that of already known members, can automatically be regarded as members. Thus, this theory allows for a finite memory representation for concepts. It captures the fact that we can determine membership of infinitely many new instances under the concepts we are familiar with, on the basis of a finite set of known facts concerning dimensions and members.

The typicality effects in nouns are robust and pervasive: *”whenever a task requires someone to relate an item to a concept, the item's typicality influences performance”* (Murphy 2002: 24). Thus, the prediction in (3a) is robustly supported by experimental research conducted in cognitive psychology. A similarity-based analysis has been applied to the semantic analysis of predicates by cognitive linguists such as Lakoff (1987), but to the best of my knowledge these applications are insensitive to word

class distinctions and their relation to the rule vs. similarity distinction. Hence, it is mainly the second type of prediction regarding adjectives as rule-based in (3b) that needs to be further explicated and justified. To this, as well as to additional problems with the gradable semantic analysis of nouns, we now turn.

2.4 Rule-based (Boolean) classification in adjectives

2.4.1 Conceptual vs. morphological gradability

In 2.3 we have presented and justified a gradable analysis of nouns. However, this type of analysis of nouns and noun phrases is often rejected by linguists, for the following reasons. Although numerous psychological facts show that, conceptually, nouns are gradable, a variety of semantic facts show that nouns are usually incompatible with degree structures. The term *gradable* is used in semantics and morphology (cf. Kennedy 1999, Baker 2003 and references therein) to refer to adjectives like *tall, bald, old, large, good, healthy,* and *clever*; these adjectives are distinguished from nouns by several characteristics. First and foremost, gradable predicates can felicitously combine with comparative and equative morphemes (as in *Sam is cleverer than Dan* and *Sam is as clever as Dan*), superlatives (as in *the cleverest*), degree modifiers (as in *very clever* and *too clever*), etc. A few adjectives, for instance, *extinct, even, married* and *nuclear,* and all nouns in languages like English (e.g., *bird, apple* and *chair*), are classified as *non-gradable*, because they cannot felicitously combine with these degree-morphemes, as demonstrated in (5).

5) a.* Tweety is more (a) bird than Tan is.
   b. *Tweet is as (a) bird as Tan is
   c.* Tweety is the birdest.
   d.* Tweety is very / too (a) bird.

The data in (5) is rather robust cross linguistically (Baker 2003). Apparent counterexamples to (5a) are normally not truly within-noun comparisons – comparisons of two entities on one noun; rather, they belong to different constructions. These constructions include between-noun comparisons – comparisons of one or two entities with respect to two different nouns, as for instance in *John is more woman than man; This vehicle is more a car than a truck*, etc. Between-predicate comparisons are discussed later on in this paper. Other constructions include comparisons whereby a noun is licensed by virtue of a mediator, like the adjective *typical of* or the preposition *of*, as in *John is less of a cook than Johanna*, These examples are different again, and are discussed separately. The crucial point is that, unless modified, nouns cannot occur in within noun comparisons and equatives, a fact reoccurring in languages as different as English, Spanish, Russian and Hebrew. In addition, nouns are utterly odd with most if not all other morphemes that reference gradability in the adjectival domain, including *too, enough, most, very*, etc. See Baker (2003: Chap. 4), for a cross linguistic justification of these generalizations.

To account for these facts, generally, semantic theories assume that in gradable predicates, unlike non-gradable ones, entities possess the properties to different extents. Entities are judged to be instances of gradable predicates iff the extent to which they satisfy the relevant gradable property is within the norm, s(P,c), as stated in (6a). The qualities or attributes that permit grading are denoted by comparative relations like *taller* or *less tall*, as in (6b) (von Stechow 1984; Klein 1991).[6]

6) a. A standard-based membership criterion:

$[[P]]_c = \{ x \in D_x \mid \deg(x,P,c) \geq s(P,c) \}.$

b. A degree-based ordering criterion:

$[[\text{more } P]]_c = \{ \langle x_1,x_2 \rangle \in D_x^2 \mid (\deg(x_1,P,c) - \deg(x_2,P,c)) > 0 \}.$

These semantic theories assume that the interpretation of non-gradable predicates, including virtually all nouns, does not involve any mapping of entities to degrees. They are directly associated with object sets, not degree functions. This explains the incompatibility of nouns with comparatives, superlatives, and other degree modifiers.[7]

This gradability approach to the noun-adjective distinction is disadvantageous in several respects. First, gradability is by no means unique to adjectives, as the discussion above illustrated. Second, gradability is not compulsory in the adjectival domain, for example, *prime* and *extinct* are not morphologically gradable. Third, a gradability approach postulates the existence of a sharp gap between conceptual and linguistic structures in the case of nouns. A significant advantage of the new thesis presented in this paper is that it suggests a principled explanation for both the psychological and linguistic facts.

On this proposal, the operations that glue together the dimensions of adjectives and nouns, respectively, are fundamentally different, resulting in rule vs. similarity categories (Sassoon 2007). The distinction between rule-base and similarity-based categorization has consequences for the type of dimensions and degree functions of adjectives and nouns. First, in gradable adjectives, the notion 'dimension' relates to gradable properties like ‘length’ or ‘blood-pressure’, while in psychological accounts of nominal concepts it also covers ordinary properties like ‘flying’. This is a direct consequence of the fact that gradability in nouns does not hinge on gradability in the dimensions; rather, it results from averaging.

Second, averaging involves normalization, namely conversion of all of the dimensional degree functions into a shared range, such as the real interval between 0 and 1. Normalization is needed because a sum/product of degrees on two different scales – for example, degrees d_1 on a scale between 0 and 1 and degrees d_2 on a scale between 6 and 100 – is not meaningful. The result of averaging is also a degree function with the standard, bounded and normalized range. The situation is very different in adjectives, which scales are often open on at least one end, and are based on conventional units or other concrete, context-driven measurements, such as meters, as in *two meters long*, number of free chairs as in *the restaurant is full/empty, except for one chair*, number of open windows, as in *open/closed except for one window*, etc. The measurements need not be normalized in ordinary usage of adjectives.

Hence, all nominal scales are alike in terms of range, whereas each adjectival scale comes with its own range. Consequently, nouns are easily comparable, as in *more a car than a truck, more a horse than a bird, more mammal than bird, more an Italian than a Japanese*, etcetera. By contrast, Kennedy (1999) has observed that most inter-adjectives comparisons are not meaningful, except in special circumstances, cf. *#longer than green*, *#taller than open*, etc. (for a discussion of conditions whereby adjectives allow such comparisons see Sassoon 2007).

Speakers treat between-noun comparisons as perfectly grammatical and interpretable, despite the fact that they cannot always say precisely what the underlying scales are that render these comparisons true or false. These facts, which are usually not noticed or not addressed by semantic theories, are hard to explain if nouns are non-gradable and not associated with ordering dimensions, but they are

straightforwardly captured by the RS hypothesis. Between-predicate comparisons are expected to favor nominal averaging functions over one-dimensional adjectival functions, because the former but not the latter are readily normalized for the purpose of averaging, i.e. they are constructed as having a shared scale, which renders them directly comparable.

Also, it is tempting to assign noun comparisons a so-called metalinguistic interpretation, e.g. *more a bird than a horse* can be analyzed as conveying "given the evidence, more justifiably called 'a bird' than 'a horse'". While this is an implication conveyed by the sentence, calling it 'metalinguistic' does not dismiss the problem of gradability in nouns. Furthermore, languages such as German and Greek have two gradability morphemes – a regular comparative equivalent to *more* and a metalinguistic *more*. Between-noun comparisons are licensed with both.

Future research should aim to test whether these comparisons are normally interpreted relative to a unique dimension or relative to multiple dimensions. Special attention should be given to between mass-noun comparisons. Examples like *more water than land* are prevalent, but they relate to measurement of quantity – the size of the area covered by water vs. dry land, rather than to degree – the extent to which some stuff exemplifies the property of being covered with water vs. the property of being a dry land. However, comparisons of more abstract nouns, as in *more show than substance*; *more rhetoric than reality*, are more plausibly between-noun degree comparisons.

Third, as discussed in section 1.2, dimensional weights are context relative. Dimensional weights may vary between people and contexts in so far as the similarity degrees they produce preserve the ordering of entities, which is crucial to capture categorization. Hence, similarity scales – i.e., averaging functions – encode a rank order of entities. By contrast, adjectival scales encode more than a mere ordering (Sasoon 2007: chap. 9; Sassoon 2010). Degree differences are meaningful. This is an important aspect, because degree morphemes, which select adjectives only, are sensitive to degree differences. Comparative morphemes directly encode degree differences, as in *two degrees warmer*, *much happier*, and *slightly angrier* – x is more than y in these cases iff the difference between x and y's degrees is two degrees, much, and slight, respectively. Other morphemes involve implicit comparison (Klein 1980), i.e. sensitivity to degree differences, including, for one, the typical adjective-selecting degree morpheme *very*, and its equivalents across languages (Baker 2003: Chap. 4). Yet, other morphemes encode degree differences, as well as ratios, as in *two meters tall*, which is equivalent to "twice as tall as a meter", *half full*, *twice as good*, and so on and so forth.

Fourth, despite the fact that both nouns and adjectives are conceptually gradable in the sense of being related with a degree function along dimensions, according to the RS hypothesis, this graded structure is implicit in nouns and explicit in adjectives. Upon processing of a nominal concept, speakers are only aware of the object set eventually construed based on each entity's averaged degree in the dimensions. Thus, speakers reject combinations such as **more bird* claiming that all birds are equally good category members. In opposition, while processing adjectival concepts speakers are aware of the graded structure and the scale it imposes, thus accepting modification of adjectives by degree morphemes.

The between-noun comparison construction is special in that it accesses nominal gradability directly, with no mediation of an adjective or any other morpheme. I hypothesize that this happens because categorization in nouns depends on similarity to a prototype, and is affected by contrast categories. In the case of two or more

contrasting categories (e.g. birds, mammals, and reptiles), entities are classified under the category to which ideal they resemble most. A between-category comparison align with this categorization strategy directly. In support of this account, notice that between-adjective comparisons improve as well, when contrasting categories within a uniform implicit dimension space are involved, as in *more red than green* and *more salty than sweet.*

Fifth, the ordering imposed by a nominal concept can also be accessed easily through the use of, for example, an adjective that takes a nominal argument and selects its dimensions, as in *better example of a bird, more typical of a bird, more normal, characteristic, representative,* or even simply *more of a bird*, which can be understood as equivalent to any one of these relations, cf. (7a) vs. (7b). As these relations are based on adjectives, whose interpretation is by default processed explicitly, a graded structure is consciously available in virtue of which the combination with a comparison morpheme is judged well-formed and perfectly interpretable. English *of* can directly access the typicality dimensions of a noun, perhaps by virtue of an elided occurrence of *typical*. This possibility is not available in all languages, cf. Hebrew. An interesting correlate is the reduced acceptability and frequency of use of *more typical of a bird*, compared to the perfectly acceptable *more typical bird*. Notice also that the latter is equivalent to "bird AND typical of a bird", meaning that, unlike the former, it entails birdhood. This renders it more appropriate in most contexts of use, since, normally, birds are under discussion when speakers explicitly relate to typicality along bird dimensions.

7) a. A robin is a better example of a bird than an ostrich.
   b. A robin is more (typical) of a bird than an ostrich.
   c. The noun activity is 'nounier' / less 'nouny' than the noun bird (Ross 1973).

Moreover, nouns turn easily into adjectives, for instance, by adding a morpheme like *-y*, as in *birdy*. The resulting adjectives are readily gradable, as demonstrated in (7c). This fact is hard to explain if nouns are directly associated with an object set, but it is captured easily by the assumption that nouns readily provide a set of dimensions. Normally, these dimensions are processed implicitly, and are bound by averaging operations. They can be processed explicitly when the noun saturates the internal argument of an adjective like *typical of* (its 'dimension-set' argument) or a preposition like *of* in (7b), or when the noun combines with an adjectival affix like *-y*. Thus, the RS hypothesis explains why it is so easy to recategorize nouns as adjectives and vice versa, and the prevalence of noun-adjective homonyms, as in *Italian* vs. *an Italian*, *poor/rich* vs. *the poor/the rich,* etc. Also, it correctly predicts the similarity and differences in interpretation in minimal pairs such as *bird-birdy*, *bird-normal bird*, *bird-typical of a bird* and *an Italian-Italian*, as discussed in 1.4.2.

Other examples of adjectives that are used to access nominal graded structures include many forms and derivatives of size adjectives as, for instance, in *a huge idiot* (see Morzycki 2009 for discussion) and quantity adjectives like *much* as in *this is pretty much a chair*; for references and review of evidence supporting analyses of *much* as an adjective see Solt (2009). Also, the WH-pronoun in exclamative constructions triggers reference to some sort of degrees, as in *what a bird!* However, the only degree construction in English in which nouns occur with no mediating adjective or other morpheme, and directly reference their degree function, is the between-predicate comparison.

In conclusion, the RS proposal is considerably improved in terms of its psychological adequacy, and its fruitfulness as a basis for explaining semantic contrasts between nouns and adjectives. It explains why nouns are conceptually gradable, yet do not directly combine with most degree morphemes. All in all, nouns behave very much like our standard formal semantics for gradable adjectives would expect. They map entities to degrees and they are linked with ordering dimensions. Nouns also exhibit vagueness effects, including borderline cases and instantiations of the Sorites paradox (see van deemter 2010), which typically accompany gradable, rather than sharp concepts (Kamp 1975). Thus, by assuming that the semantic interpretation of nouns is non-gradable, linguists pay a heavy price in terms of the dissociation between the semantics they assume for nouns and many other things that we know about them. All things considered, the assumption that nouns *are* semantically gradable seems better off.

In addition, nouns combine with *more-of* comparatives, and with degree modifiers like *pretty much*, and they are licensed in between-predicate comparisons of the form "more P than Q" more freely than adjectives are. The infelicity of nouns in, for instance, within-predicate comparisons, must have reasons other than lack of gradable meaning; a possible reason is an underlying similarity-structure, which yields ordinal and implicit gradability.

1.4.2 Conceptual dimensions vs. linguistic 'respects'

Semanticists assume that the interpretation of one-dimensional adjectives like *tall* is based on a single dimension, e.g., 'height' (Kennedy 1999), whereas the interpretation of multidimensional adjectives is based on several dimensions (Kamp 1975); e.g., the adjective *healthy* orders entities along a number of dimensions, such as *blood pressure, pulse,* and *lung functions,* or *flu, pneumonia* and *chickenpox*. The adjective *intelligent* can be ordered on dimensions such as *mathematics, literature* and *personal relations*, etc. Adjectives like *human, childish, good, nice, beautiful, optimistic, positive* and so on, can relate to multiple physical and/or behavioral respects. The range of dimensions in the interpretation of adjectives is highly context dependent (Klein 1980: 6-8).

Significantly, the contextually relevant dimensions of a multi-dimensional adjective can be overtly specified as part of its argument structure, using a *with-respect-to* (henceforth – wrt) prepositional phrase, as in *healthy wrt blood pressure.* In addition, grammatical operations can access the dimensions of multi-dimensional adjectives and operate on them (Bartsch 1986; Landman 1989). For example, we can quantify over these dimensions or respects, as in *healthy in every respect* and *generally healthy*. The oddness of combinations like *tall wrt height* and *tall in every respect* is likely due to the fact that *tall* is a one-dimensional adjective. As for nouns, although they are clearly associated with ordering dimensions, for reasons which are not immediately obvious, they differ from gradable adjectives in that their dimensions cannot be accessed by grammatical operations such as those denoted by wrt-phrases and quantifiers, as demonstrated in (8).

8) a. Tweety is healthy in every respect.
   b. Tweety is generally healthy.
9) a. #Tweety is a bird with respect to flying / size.
   b. #Tweety is generally a bird.

Hence, noun dimensions do not typically function as arguments. We do not tend to describe an entity as *a bird in every respect except for flying,* or *a chair except in the seat sense* or *except for not being a seat*. Such statements are not flat-out ungrammatical, but speakers generally agree that they are marked and unusual, cf. the contrast with similar adjectival statements in (8) vs. (9). Exceptions to this generalization include nouns that are directly related to an adjective semantically and usually also morphologically, for example, nominalizations such as *similarity wrt color*, and evaluative nouns such as *a genius/idiot except with respect to mathematics*. Those are discussed in part 3. Additional notes on different types of usage of wrt-phrases are found in the appendix.

These facts are explained by the RS hypothesis. Intuitively, modifying a predicate P with a wrt-phrase makes sense only when several dimensions are treated as necessary conditions for membership in either its positive or its negative denotation, and as a consequence, entities may indeed be regarded as P in one respect and not P in another. Thus, a predicate P can be modified by a wrt-phrase, or can assign a 'wrt' argument-role, iff each of P's dimensions can function as a categorization criterion, as in *healthy wrt blood pressure, but not healthy wrt lung functions*. Conversely, in nominal concepts like *bird* or *not a bird*, the dimensions are normally not necessary for membership. At best, some of them are very important. Thus, nouns normally do not license 'wrt' phrases. Hence, the RS hypothesis explains the licensing of wrt-phrases by virtue of the association of adjectives and nouns with rule- vs. similarity-based dimension integration.

According to the RS hypothesis, grammatical operations such as quantifiers and wrt-phrases cannot access the nominal dimensions, because the latter are bound by an averaging operation at an early processing stage. This averaging yields immediate categorization judgments, resulting in a set of instances. Thus, the dimensions involved go unnoticed. Put differently, nominal dimensions are not handled each one separately, they are automatically integrated via averaging, which is incompatible with having a dimensional argument role.

In support of this account, notice that when nominal dimensions do add categorization criteria, a wrt-phrase is licensed. For instance, if an expert characterizes birds by the possession of, say, 100 separate genes, which all and only birds possess, she might indeed describe new species that possess only 50% of these genes, as *birds in this respect, but not in that respect*. This is clearly a scientific context in which the noun's usage is based on a definition, a rule deviating in nature from the similarity structure on which ordinary, daily interpretations are based. Rule-based processing goes hand in hand with the licensing of a respect argument. [8,9,10]

In addition, in the dimension set of one-dimensional adjectives like *tall*, we cannot find two different dimensions to form categorization conditions for *tall*, so the requirement for the licensing of a wrt-phrase is not met. All considered, a wrt-phrase can only be licensed in multi-dimensional adjectives. Like any other argument, the wrt-argument can interact with determiners, as in *is healthy in every respect*; a determiner which quantifies over respects states how many of the dimensions form necessary conditions for membership in the category in the context of use.

In sum, adjectives have a 'respect' argument. When their respect argument is saturated, as in (10c), or bound, as in (10a,b), the dimensions function as categorization criteria**,** which amounts to a rule-based analysis. Finally, wrt-modification and quantification over dimensions are possible when a noun saturates an adjectival argument as in (10e), the adjective accesses the nominal dimension set; the syntactic category 'Adjective' triggers a change in the mode of processing of the

dimensions. Instead of dimension integration by means of averaging, through the similarity system, the dimensions get integrated via logical operations, through the rule-based system. The procedure is explicit – open to introspection. The number of dimensions selected can be consciously monitored and they can be expressed using grammatical operations, as illustrated in (10).

10) a. This bird is normal/ typical in every respect.
    b. This bird is generally atypical.
    c. The boxes are identical with respect to color (but different in size).
    d. This bird is normal/typical, except for its big size.
    e. With respect to size/flying, Tweety is similar to a bird.
    f. *With respect to size/flying, Tweety is a bird.

Importantly, the RS hypothesis does not exclude the possibility of implicit processing of some parts of an adjective's interpretation; rather, it requires awareness of a certain categorization dimension. For example, colour concepts are construed based on aspects such as saturation and hue; these dimensions were revealed through empirical research; standard speakers are hardly aware of their role in colour categorization – they are processed implicitly. Thus, these dimensions do not function as adjectival respects. Speakers explicitly associate each colour term with a single scalar property which they regard as a basis for categorization, i.e., colour predicates are conceived by speakers as one-dimensional. These single scalar properties, then, count as respects. Indeed, across languages, color terms are usually adjectives (Dixon 1982) and infants acquire them significantly later than nouns (Bornstein 1985).[11]

An additional piece of evidence supporting a different analysis of nominal and adjectival dimensions is, then, the fact that natural languages refer to them by different names. The adjectival dimensions are called *respects*, as in *Dan is not healthy in three respects: bp, pulse and sugar*. By contrast, the nominal dimensions are related to as *typical* or *characteristic*, as in *flying, singing and perching is typical of birds*. This would be unexpected had these dimensions been exactly the same. In fact, the adjective *typical* has two related interpretations. On one interpretation, it predicates over entities, as in *Tweety is a typical bird*, whereas on another, over dimensions, as in *Eating fish is typical of water birds*. Here, the subject NP *eating fish* does not reference any birds; rather, it designates a dimension of birds. While I use *typical* as the running example here, *normal* is another alternative; the difference between the two is a topic I leave for the future. Likewise, we can say that *normally, birds fly and eat insects*, but saying that *normally, healthy is healthy with respect to blood pressure* is utterly impossible.

In conclusion, psychological theories that treat nouns as gradable and multi-dimensional, fail to explain important semantic contrasts. A main difficulty with incorporating a gradable analysis of nouns (and noun phrases) into the semantic theory is that important distinctions might become blurred. First, if nominal concepts denote degree functions, it is not clear why they are incompatible with within-predicate comparatives, equatives and other degree modifiers. Second, if nominal concepts are multi-dimensional, it is not clear why it is impossible to quantify over their dimensions. A larger problem is looming behind these questions, namely the problem of giving an adequate account of the semantic distinction between nouns and adjectives. A major advantage of the RS hypothesis is that it suggests a principled explanation for these facts, which hinges upon distinctions in the type of graded

structure (degree function and dimension type) and underlying processing of nominal and adjectival concepts.

1.4.3 Bare multidimensional adjectives as rule based

When adjectival dimension arguments are explicitly saturated, as in *healthy wrt blood pressure and lung functions*, or bound, as in *healthy in every respect*, they function as categorization criteria. However, an open question is what happens when a multidimensional adjective occurs bare, as in, e.g., *John is healthy*. Do the dimensions integrate through logical operations such as conjunctions and quantifiers also in this case? A-priori, it makes intuitive sense to think that in this case we use averaging. However, judgments of speakers with respect to this question reveal a different picture (for survey results see Sassoon 2011b).

Consider, for example, a context in which health is measured by the dimensions *blood pressure, pulse* and *sugar* (a measure of diabetics). Imagine that Dan has the maximal degree in two of these dimensions, but is not within the norm in the third. Conversely, imagine that in all these dimensions, Sam's levels are within the normative range, but they are the lowest possible, so Dan's average score on the dimensions is higher. Intuitively, in this scenario, Sam is strictly speaking healthy, but Dan is not, because Sam, but not Dan, reaches the norm in all the contextually relevant respects. Because of that, intuitively, Sam is healthier than Dan. Hence, it is not the case that we directly compare Sam's and Dan's averages on the dimensions. Had we done that, we would have judged Dan to be healthier than Sam. Rather, first, we fix negative and positive denotations for healthy, based on dimension conjunction; for the positive denotation, we select entities that reach the standard in all the dimensions. Afterward, we fix the ordering relation to be such that positive denotation members are always healthier than negative denotation members.

Consider also a context whereby dimensions are weighed. Assume Dan is an athlete with ideal levels of blood pressure and pulse, but he has asthma, i.e. occasional attacks, for which he has the necessary medication. Sam, again, has no asthma, but, being very sedentary, has higher blood pressure (13/8) and pulse (78 bpm). These values are still within normative limits. Would we still call Sam healthier than Dan? The answer is twofold. First and foremost, strictly speaking – yes! If Sam is healthy in all respects, despite not being an athlete, while Dan is asthmatic, Sam is healthier than Dan. Second, pragmatics plays a significant role in the interpretation of adjectives. Speakers tend to enrich interpretations in various ways. Thus, if informed that Sam is very sedentary, they may infer that Sam is not healthy in ALL respects. In that case, Sam is predictably not healthy. In addition, recall that domains of operations referenced by natural language expressions are virtually always restricted to contextually relevant entities. Thus, Dan may be considered healthier than Sam in contexts in which asthma is ignored.

If these observations are on the right track, they support a systematic semantic difference between multidimensional nouns and adjectives. While nominal denotations consist of entities whose average on the dimensions is sufficiently high, denotations of adjectives like *healthy* and *typical* consist of entities that fall under *all the dimensions*. Hence, a person, healthy in every respect except the flu, strictly speaking, is *not healthy*. By contrast, a creature satisfying all features characteristic of birds except that it doesn't fly or lay eggs, may still be a *bird*, albeit not a *typical* one. Thus, the default interpretation of adjectives like *healthy* involves universal quantification over dimensions. Let us call these adjectives 'conjunctive', as stated in

(11b). Furthermore, the denotation of adjectives like *sick* and *atypical* consists of entities that have some disease or other, possibly only one. Thus, the default interpretation of these adjectives involves existential quantification over dimensions. Let us call these adjectives 'disjunctive', as stated in (11c).

11) Let $F(P,c) = \{F_1 \ldots F_n\}$ be the dimension set of P in context c; $\forall x \in D_x$:
   a. If P is *nominal*, $x \in [[P]]_c$ iff $deg(x,P,c) \geq s(P,c)$ (cf., 1.2).
   b. If P is a *conjunctive-adjective*, $x \in [[P]]_c$ iff $\forall F \in F(P,c)$: $x \in [[F]]_c$,
      iff $x \in [[F_1]]_c$ & … & $x \in [[F_n]]_c$.
   c. If P is a *disjunctive-adjective*, $x \in [[P]]_c$ iff $\exists F \in F(P,c)$: $x \in [[F]]_c$
      iff $x \in [[F_1]]_c$ or … or $x \in [[F_n]]_c$.

Typical and *normal* are conjunctive multi-dimensional adjectives. When they modify a noun P, the resulting interpretation is mediated by a universal quantifier over the nouns dimensions, as demonstrated in (12).[12] This proposal derives the intuition that the adjectival phrase "typical of P" is stronger than the nominal predicate P – it has more categorization criteria – although it is hard to put a finger on the exact dimensions which add criteria, since the dimension set is context dependent, and since it is always possible that interpretation involves an implicit 'wrt' modification, as in *typical wrt flying*. The dominance of dimensions like *feathered* to the concept *bird* is lost when considering the respective adjective *typical of a bird* (Hampton 1979, 1998); the RS hypothesis predicts this, since additional, less important *bird* dimensions turn into categorization criteria of *typical of a bird* along with *feathered*.

12) a. $[[\text{Tweety is a typical bird}]]_c = 1$ iff $[[\text{Tweety}]]_c \in [[\text{bird}]]_c$ &
   $\forall F \in F(\text{typical of a bird},c)$, $[[\text{Tweety}]]_c \in [[F]]_c$
   (Tweet is a bird and is typical of a bird in every respect).
   b. $F(\text{typical of a bird},c) = \{\text{typical wrt size, typical wrt flying,}\ldots\}$
   c. $[[\text{Tweety is typical of a bird wrt size}]]_c = 1$ iff
   $| deg([[\text{Tweety}]]_c,\text{size}) - deg(\text{bird},\text{size}) | < n$

The context-dependency of the adjectival dimensions makes it hard to experimentally support or refute a conjunctive dimension-integration hypothesis. It is hard to say which and how many dimensions count in each context. However, a new corpus-based method can be used to overcome these difficulties, which is based on measurement of the frequency of markers of universal quantification – exception phrases – with adjectives and nouns. The licensing of exception phrases depends on mediation of universal or quasi universal quantification, as illustrated in (13a-b).

13) a. **Everyone** is happy **except** for Dan.
   b. **No one** is happy **except** for Dan.
   c. #**Someone** is happy **except** for Dan.

(13c), is, therefore, infelicitous, except in an alternative non salient "in addition to" interpretation; the interpretation that interests us here is unavailable, namely the one whereby Dan presumably is not happy, as in (13a).

Exception phrases can operate on dimension sets, as in *healthy except for high blood pressure*. Thus, compatibility with exception phrases can form evidence for the hypothesis that the interpretation of an adjective involves universal quantification over dimensions, rather than existential quantification or averaging. Our proposal

predicts that *except* will be freely licensed as an operation on the dimension-set of some adjectives – the conjunctive ones. This prediction is borne out, as the naturally occurring examples in (14) illustrate.

14) a. I am a 64-year-old man, quite healthy except for high blood pressure...
b. Sam's early development was considered typical except for slight articulation errors noted in kindergarten which resolved spontaneously.
c. … my Mother's family, mainly tradish, eat more of the tradish foods and they seem to be healthier except the cancer aspect.
d. $[[\text{Dan is healthy except wrt blood pressure}]]_c = 1$ iff
$\forall F \in (F(\text{healthy},c) - \{(\text{healthy wrt}) \text{ blood pressure}\}): [[\text{Dan is F}]]_c = 1$
(Dan is healthy wrt every dimension except blood pressure).

Conversely, the RS proposal predicts that this will rarely occur with disjunctive-adjectives like *sick*, whose default interpretation is mediated by an existential quantifier. A non-default universal interpretation is likely to be explicitly marked, as in *sick in every respect*. Example (15) shows that indeed *except* cannot operate on the dimension set of a bare disjunctive adjective, and in fact, Google searching for key-words like *sick except* and *atypical except* barely provides any examples with exception phrases operating on an implicit universal quantifier over dimensions, in contrast with abundant examples it provides for *healthy* and *typical*.

15) #They are sick, except for (normative) blood pressure.

Finally, on this proposal a *negated* disjunctive adjective like *not sick* denotes entities that have *no* disease. Likewise, (16a) is an example of a negated use of *atypical*, which conveys that the children were atypical in *no* respect, except for their intelligence. Crucially, *no* is a universal quantifier. Thus, we predict that exception phrases will occur operating on the dimensions of negated disjunctive adjectives; this turns out to be the case, as briefly illustrated in (16).

16) a. Apparently, the children were not at all atypical, except that they were brighter than the average high- school Senior.
b. They do not appear to be sick, except for the diarrhea.
c. $[[\text{Dan is not sick except wrt blood pressure}]]_c = 1$ iff
$\neg\exists F \in (F(\text{sick},c) - \{(\text{sick wrt}) \text{ blood pressure }\}): [[\text{Dan is F}]]_c = 1$ iff:
$\forall F \in (F(\text{sick},c) - \{(\text{sick wrt}) \text{ blood pressure }\}): [[\text{Dan is F}]]_c \neq 1$
(Dan is sick wrt no feature, except blood pressure).

Generally, negated universals are existential, and vice versa, as the equivalences in (17a-b) illustrate; hence, exception phrases are expected to combine with negated-disjunctive, but not -conjunctive adjectives, as the contrast in (17c-d) illustrates.[13]

17) a. (healthy ⇔ healthy in *all* respects) **iff** (not-healthy ⇔ not-healthy in *some* respect).
b. (sick ⇔ sick in *some* respect) **iff** (not-sick ⇔ sick in *no* respect).
c. Dan is not sick, except for blood pressure (= sick in no respect except b.p.)
d. ?? Dan is not healthy, except for high blood pressure.

Sassoon (2011a) tested these predictions quantitatively The data found in balanced linguistic corpora of English such as the BNC and COCA is consistent with the above predictions, but is very scarce because of the low frequency of negated forms and exception phrases. Hence, data was drawn from the internet using the search engine Google. The study did not rely on Google's estimations of number of hits, which are often unreliable. Moreover, the examples for this study were scanned one by one, and repeated entries and entries clearly exhibiting deficient English competence were removed. Lapata and Keller (2005) show in a variety of ways that, despite the many hits by non-native speakers, Google-based counts correlate with frequencies obtained from a carefully edited, balanced corpus such as the BNC, and they reliably predict Native-English judgments.

The sample consisted of 1814 naturally occurring examples of exception phrases preceding a predicate. All the items searched for were put in inverted commas (as in: "healthy except"). Before the removal of repeated and non-English entries, the sample of exception phrases preceded by an adjective consisted of the first set of up to 100 Google-search results with each of 18 adjectives. After removal, the sample consisted of a total of 1444 examples (M = 80). The adjective list included *normal, typical, healthy, familiar, healthier, bad, sick, atypical, abnormal, different, identical, similar, good, better, intelligent, dissimilar, worse,* and *unfamiliar*. In addition, after removal of repeated/non-English entries, the sample of exception phrases preceded by a noun consisted of a total of 368 counts, with five nouns. This sample included the first 100 Google-search results for the noun probes *bird except, table except* and *mother except*, as well as the first 34 Google-search results for each of two additional noun probes *capital except* and *carrot except.* The nouns *bird* and *table* are typical count nouns, whereas the other three nouns are not paradigmatic category-denoting nouns. *Carrot* can function as a mass-noun, as in *there is much carrot in the salad; Mother* is +human, and *capital* has an adjectival homonym. In addition, in its nominal use, as in *the capital of France*, it is relational, and so is the noun *mother*; i.e., these nouns exhibit argument structure, a property more typical of adjectives than of nouns. Given these similarities to adjectives, one could expect these nouns to resemble adjectives also with respect to the licensing of a dimension argument, and in accordance, one could expect to observe exception phrases operating on their dimensions (henceforth, 'dimension set readings'), as in, for example, *mother, except biologically*.

The counts were classified into 4 categories. First, counts were classified as either *negative* or *positive* depending on whether the adjective did or did not occur in the scope of negation. Second, they were classified as either 'dimension-set readings' or not. Briefly, exception phrases that did not operate on a quantifier over dimensions, but rather over entities, events, cases, times, degrees, parts, etcetera, were classified as 'non-dimension set readings', as in *Everyone's been sick, except me,* which exemplifies quantification over entities, and *The tests are identical, except for one question*, which exemplifies quantification over parts. Exception phrases were similarly classified as 'non-dimension set readings', if an explicit (quasi) universal expression was present, like *everything, nothing, little, most, mostly, much, totally, completely, absolutely, otherwise, never,* and *all in all*, as in, for example, *Nothing abnormal except for high BP*; an expression like *nothing* suffices to license an exception phrase; thus, we cannot use this example as evidence that *healthy* triggers accommodation of an implicit universal quantifier. Additionally, exception phrases related to a different clause, not to the adjective preceding it, were also classified as 'non-dimension set readings', as in *You would never know I was sick. Except for being bald, I look great.*

Based on these classifications, the frequency of dimension-set readings, i.e., exception phrases operating on an implicit quantifier over dimensions, was calculated for positive and negative contexts.

The results were significantly different for the adjectives and the nouns. More than a third (36.5%) of the 1444 exception phrases preceded by an adjective in the sample were of the form "adj. except a dimension", involving an implicit quantifier over dimensions. These examples occurred with all the adjectives. By contrast, none (0%) of the ~370 "noun except" counts in the sample were of the type "noun except a dim". Hence, **hundreds** of dimension-set uses were found with the tested adjectives, confirming logical, rule-based dimension-integration, while **no** such uses were found with the nouns, supporting an averaging-based integration of their dimensions. A finer grained analysis of the results in positive vs. negative contexts supports a two-way adjectival typology, with predominantly conjunctive and predominantly disjunctive adjectives, in accordance with the thesis of this paper, e.g., the predictions (14)-(17) illustrate. Additionally, an unexpected generalization emerged from the data: The positive adjectives were all predominantly conjunctive and the negative ones all predominantly disjunctive (Sassoon 2011a; for similar results on surveys of acceptability judgments see Sassoon 2011b). Finally, moderate to high correlation was found between conjunctivity, as measured by the frequency of dimension set readings in positive contexts, and the frequency of occurrence of an adjective with *perfectly*, as estimated by Google.

In sum, these findings give preliminary support to the proposal that the dimensions of adjectives integrate by means of logical operations, not similarity functions; multidimensional adjectives sharply differ from nouns in this respect, and they resemble artificially construed rule-based categories. Furthermore, antonymy and licensing of *perfectly* appear to be predictive factors of conjunctive vs, disjunctive rule-based categorization. However, since the no. of lexical items tested is small, the generality of this proposal should be tested on additional nouns and adjectives.

This brings us to the end of part 1 of this study. We discussed a variety of arguments to the effect that similarity-based processing is the natural option for nouns, while rule-based processing is an available and preferable option for adjectives. We now turn to the acquisition and neural basis of nouns and adjectives.

2 Neural and developmental predictions of the RS hypothesis, existing support

The rule vs. similarity distinction has well-studied neural correlates and corresponding developmental correlates. This section compares them to developmental and neural findings from the literature about adjectives and nouns. An emerging analogy would yield indirect support to the RS hypothesis. An alternative approach whereby categorization in both nouns and adjectives is based on the same type of structures, would predict no differences in processing, and accordingly, in the neural basis and developmental course of nouns and adjectives.

Rule-based classification tasks involve explicit hypothesis testing procedures. Explicit reasoning demands working memory and executive functions, thus recruits primarily frontal-striatal regions (including the anterior cingulate, the prefrontal cortex and the head of the caduate nucleus). Conversely, similarity tasks involve implicit procedural memory; thus it recruits regions such as the tail of the caduate nucleus, which do not project directly to the prefrontal cortex. This corresponds with non awareness of the similarity-based categorization rule (for the full neuropsychological model and its motivation see Ashby and Maddox 2005: 163-165).

The neural dissociation between rules and similarity-based categories has developmental consequences. Children younger than 3 years old consistently perform similarity-based processing tasks. Conversely, due to the late maturation of the prefrontal cortex, by age 3 children tend to still have difficulties in consistently using rules, which only get resolved at age 5 (Frye et al. 1995; Zelazo et al. 1996, 2004; Thomason 1994). Furthermore, children up to age 10 often base categorization on similarity even where definitional properties exist (Keil 1989).

For example, numerous studies (Frye et al. 1995; Zelazo et al. 1996) show that given a pile of cards with red and blue triangles, 3 year olds succeed in tasks such as "if red, put here, otherwise put there"; yet, they fail when circles are added, and the task demands paying selective attention to one of two dimensions – color or form. After responding correctly to a task such as "if red, put here, otherwise put there", children fail to perform correctly with a new rule such as "if triangle, put here, otherwise put there", as this task requires suppressing the old rule and attending to a new one. Also, children fail when the task demands attention to a conjunction of dimensions, as in, for instance, "if a red triangle, put here, otherwise put there". These results were obtained with different types of stimuli. In addition, performance improves with age, adults doing better than children; but as adults grow old, their performance deteriorates again.

Thus, selective attention to one of several dimensions or to dimension conjunction or disjunction is demanding. According to the RS hypothesis, this is precisely the capacity required for the interpretation of adjectives.

First, consider one dimensional adjectives. It is difficult to focus on a single dimension, because it is difficult to suppress 'irrelevant' dimensions. To understand this, it is useful to consider one important bias of children in acquisition. It is widely held that infants profit from lexical biases, which reduce the set of potential meanings of novel words. One of these biases is *the whole object constraint* – children are biased to assume that novel words like, for instance *dog*, refer to a whole object – a dog – not to an attribute or to a part of a dog (Markman, E. M., & Jaswal, V. K. 2004). To see how this constraint is related to the present discussion, consider, for example, a red pencil. Intuitively, pointing to such an object and uttering a novel word such as *goobar* strongly suggests that *goobar* means pencil, not *red*, *wooden*, or any other prominent quality of the object. It is this bias that the whole object constraint dictates.

Why are objects preferred to qualities (dimensions), rather than vice versa? A plausible answer is that there are many dimensions. Taking all of them into account provides an easier resolution of reference than trying to decide which one counts as relevant. To learn and use names of dimensions we need to suppress all the other dimensions of objects. Moreover, it is easier to imagine an object without one of its subparts, e.g., me without my left hand, than to stripe a dimension out of an object. It is difficult to think about me without shape, size or color, as the result is not an object at all. Therefore, processing of single dimensions is, arguably, more difficult than processing parts of objects, as the latter are objects on their own right.

Second, consider adjectives with multiple dimensions, which have a dimension argument. Recall that this argument can be saturated (as in *healthy with respect to blood pressure*; *talented/ good in mathematics*), or bound (as in *healthy in every respect, generally sick, no different except for size*). When an adjective occurs bare (as in *Dan is healthy / sick / different*), the dimension argument may be implicitly saturated or bound. This means that speakers and listeners have to pay selective attention to one of

multiple dimensions, or to conjunctions or disjunctions of dimensions, suggesting that adjectival processing is demanding in the same way rules are.

In opposition, nouns like *bird* do not normally license a dimension argument, or quantification over dimensions. Hence, noun processing does not require selective attention to one of several dimensions, nor to conjunctions or disjunctions of dimensions. It should therefore be less demanding in terms of working memory and executive functions. Let us consider developmental studies that indirectly test these predictions.

In terms of acquisition order, acquisition of adjectives – i.e., relatively frequent and consistently correct use – is significantly delayed compared to nouns (Damon et al. 2006: 315), as illustrated shortly. Nouns are learnt earlier, faster and with fewer errors. This is surprising given that nouns are typically characterized by a wide range of features, so many so as to defy definition (Wittgenstein 1953), while adjectives associate with a single dimension, or a relatively restricted set of dimensions (health measures, intelligence measures, etc.) The RS hypothesis provides a simple explanation to this puzzle. It is precisely the need to suppress irrelevant dimensions that makes the acquisition and use of adjectives more difficult. Moreover, a difference in acquisition rate between nouns and adjectives typically occurs in children 3-5 years old. The time course of acquisition, then, corresponds with the developmental facts reported concerning rule based categories (Ashby and Maddox 2005). Importantly, at age 3, children already gain control of meaning components of one-dimensional adjectives like *tall*, namely of comparison classes, standard types, etc. (Syrett, Kennedy and Lidz 2009). However, they still do not use adjectives consistently.

Several studies report that adjectives are still rare in the output at age 3, including adjectives which are most frequent in adult usage, such as color and size adjectives. This is reported for diverse languages, including, for example, English and Spanish (Gassar and Smith 1998: 269-271), Russian (Gvozdev 1961: 437-8), and Hebrew (Berman 1988). However, a recent cross-linguistic study shows that, by age three, children already reach the frequency of adjectives in child-directed speech, namely the speech caregivers direct to children (Ravid et al 2010).

Still, Ravid et al (2010) argue that the acquisition of adjectives lags behind that of nouns, and to depend on it. Authors generally hold that adult-like production frequencies do not always reflect adult-like understanding; in the case of adjectives, the latter lags behind the former. Errors in understanding are frequent with adjectives at age 3-5, consisting, among other things, in confusion between dimensions, children replace, for example, *dark* with *loud* or *tall* with *wide*; also, children confuse between antonyms, using, for instance, *tall* instead of *short*; in response to questions such as, e.g., *what color is it?* children may confuse *red* with *green* (Gassar and Smith 1998).

Evidence for the late acquisition of adjectives is manifested also in cases of incomplete acquisition. Polinsky (2005) has compared between 4 competent speakers of Russian and 5 adult incomplete learners, whose acquisition of Russian was interrupted at age 4 (n=3), 5 (n=1), and 6 (n=1). The incomplete learners turn out to perform significantly better with nouns than with adjectives (Polinsky 2005). For example, recognition time is longer for adjectives than for nouns only in incomplete learners. Also, translation of adjectives to a second language is less accurate, and, unlike translation of nouns, it is based on words from other classes. These word class effects are stronger than frequency effects (Polinsky 2005: 423). A natural explanation to them is that the use of adjectives requires selective attention to one or more of multiple dimensions, which is difficult before age 5, and sometimes long afterwards.

Similar results are obtained in studies of learning of invented nouns and adjectives. One and a half year olds learn nominal labels of invented categories ("this is a dax") efficiently and remember them well over several days and weeks, while three year olds still have difficulty learning adjectival labels ("this is a dax one"; "dax, not red"). Learning of adjectival labels has proved modest at best (Waxman and Lidz 2006). Moreover, a strong linkage between nouns and object categories (i.e. similarity-based categories) occurs early on, while adjectives are linked to properties (rule-based categories) only at a second developmental stage.

For example, Waxman and Kosowski (1990) show that, at 12-13 months, novel words presented either as count nouns or as adjectives direct infant attention to object categories. Thus, the property-category distinction between count nouns and adjectives does not appear in the initial phases of word learning. This is compatible with the proposal whereby, due to the late maturation of the prefrontal cortex, children tend to process rule based categories as similarity based ones.

In Booth and Waxman's (2003) study, slightly older, 14-month-old infants have shown sensitivity to the semantic word class distinction, but still failed to associate adjectives with properties systematically. They were introduced to four toy objects sharing a category and a property value (for example, purple horses). The experimenter labeled these objects either with novel nouns (e.g., "This one is a blicket"; ''These are blickets'') or adjectives ("This one is blickish"; ''These are blickish''). In addition, infants viewed an object whose category and property value are different (e.g., an orange carrot), which was labeled with a negated noun (''not a blicket'') or adjective ("this one is not blickish"; "not a blickish one''). Finally, the experimenter introduced two test objects differing in either the property (e.g., a green horse) or category (a purple chair). When infants were asked to 'find another blicket' (the noun condition) they strongly favored the category match. Yet, when asked to 'find another blickish one', infants showed no preference for either test object.

These findings were obtained in different measures (word extension, novelty tasks, an automated procedure, etc) and were replicated in numerous studies (Waxman and Booth 2001, Booth and Waxman 2009 and references therein). In conclusion, by 14 months of age, infants are sensitive to the property vs. category distinction and word classes influence infant decisions as to whether the extension of a word is a similarity based category or a property (rule). Infants more consistently relate nouns to similarity-based categories than adjectives to properties (rules).

By age three, children acquiring English distinguish novel words presented as nouns from those presented as adjectives (e.g., ''This is a fopin'' vs. ''This is a fopish one''), and they interpret novel words presented as nouns as referring to categories of objects; however, they still often interpret novel adjectives as denoting object categories, too; they do so less often when the objects in question are known to have a nominal label, but even then children have difficulties (Markman and Jaswal, 2004; Waxman and Lidz 2006 and references therein).

Berman (1988) provides a different type of evidence showing that adjectival acquisition differs from that of nouns. Adjectives resemble verbs in that they typically function as predicates not arguments. However, in Hebrew, adjectives resemble nouns in that they inflect for gender, number and definiteness, but not for tense and person. Accordingly, acquisition studies in Hebrew reveal the following differences between word classes. Verbs are learnt in morphological paradigms, whose items are interchanged with each other; for example, *ligdol* 'to grow up' may be interchanged with *legadel* 'to grow (something)' or with *lehagdil* 'to enlarge'. By contrast, nouns are learnt individually; e.g., nouns like *godel* 'size', *hagdala* 'enlargment', *gidul*

'plant, tumor' and *gdula* 'virtue, high merit, talent' are not interchanged; rather, nouns are learnt with their semantic fields (for instance, spoon-fork-knife). Adjectives, in turn, are learnt later, first with their semantic contrast set, like nouns (e.g., big/small; red/blue/…, etc.), then with their morphological paradigm, like verbs, e.g. *gadol* 'big' with *megudal* 'enlarged', etc.

A possible explanation for the acquisition lag may hinge upon the fact that an adjective's interpretation depends on the noun it modifies (cf. Pustejovsky, 1991); this can be illustrated by the different interpretations of *fresh* when combined with different nouns (eggs, air, starts, ideas, etc.) Similarly, a fast typist types quickly, a fast car can move quickly, and a fast waltz has a fast tempo. Notice however, that nominal categories also exhibit a wide range of interpretations. For example, there are different bird kinds, there are literal as well as metaphoric interpretations for *bird*, etc. At the end of the day, both nouns and adjectives reference sets of objects, with the minimal difference that the categorization rule for adjectives makes use of relatively fewer dimensions, possibly single ones.

In line with the RS hypothesis, Gassar and Smith (1998) argue that nominal and adjectival concepts both denote categories, but these categories differ markedly in their size, overlap and, importantly, number of relevant dimensions. They show that these factors may affect acquisition rates and processing difficulties. Importantly, they illustrate that category size and overlap typically depend on the number of category dimensions. One can imagine the set of possible objects as a space whose axes are dimensions such as size, texture, brightness, and so on (perceptual properties). One-dimensional adjectives such as *little* and *dark* are applicable to a very large proportion of this space of possible objects; since for these words most sensory dimensions are completely irrelevant, many different kinds of objects can be dark and many can be little. Nouns such as *dog* and *box*, on the other hand, apply to a very small proportion of the space of possible objects, since their referents must fall within a restricted range of values in many different dimensions. Thus, all dogs are, in comparison with, e.g., all the little things, very much alike. For the same reasons, adjectival categories typically overlap (red things can be big or small) but concrete basic level noun categories typically do not (a dog cannot be a house or table).

Gassar and Smith (1998) propose that it is these kinds of differences that make nouns easier to learn than adjectives, demonstrating that the noun advantage could emerge through such differences alone by means of a connectionist network that learns invented 'nominal' (i.e. multidimensional similarity-based) categories faster than 'adjectival' (one-dimensional) ones. Acquisition was measured by the number of training instances required for the network to learn the categories, as well as by error types. On each run, the network was given 1000 pairs of input and output (i.e., object representations together with appropriate labels for them); adjustment of the dimensional weights in the network followed the presentation of each pair. Then the network performance with 500 new inputs was measured. The results of experiment one (Gassar and Smith 1998) consist of a close to perfect performance on the nouns by the 4th run, and a continuous improvement in performance on the adjectives, which by the 10^{th} run still do not reach the noun level.

In response to adjectives, the network models young children's difficulty in attending selectively to individual dimensions (Gibson 1969; Smith 1989). The difficulty with adjectives occurs when distracting dimensions are available, e.g., black vs. white is difficult iff attention to size, shape, etc. has to be suppressed. What is more, given a specified dimension, for example, *red*, the probability that the network responds with *orange*, *yellow*, *green*, *blue*, or *purple* is greater than the probability

that it responds with *big*, *rough*, or some other non-color adjective. The network 'knows' that *red*, *blue*, and *green* are words of the same dimensional kind before it knows which specific sensory inputs are red. This knowledge derives from early association of adjectival outputs with the appropriate linguistic-context (their explicitly taught dimension). This type of errors is characteristic of children.

In experiment 4, nominal categories were organized by an equally restricted range of variation on all four dimensions, while in adjectival categories the range of variation on some dimensions was wide and on others narrow, as is typical of adjectives. The nominal categories were learned significantly more rapidly. Experiment 2 exhibited the same effect where the crucial difference was in category size (nominal categories are typically smaller, as explained above). In experiment 3, nouns were construed one-dimensionally; the only factor distinguishing adjectival from nominal outputs was the lack of direct learning of the lexical dimensions of nouns. This factor seems to play a role in the nominal advantage, although to a lesser extent.

Finally, given the RS hypothesis, adjectival dimensions combine through logical operations. It is therefore important to see whether acquisition (consistent semantically correct use) of logical words referring to logical operations such as conjunctions and disjunctions, or universal and existential quantifiers, is in fact delayed till at least age 5. While the pragmatics of such words (derivation of scalar implicatures or failure to do so) is a vivid topic of investigation, hardly any studies examine the acquisition of the bare semantics (truth conditions) of these words. The one study I am aware of (Paltiel-Gedalyovich 2003), focusing on Hebrew, supports the predictions of the RS hypotheses.

In accordance with the thesis that the demands of logical processing are dependent on the prefrontal cortex, which maturates late, the acquisition of words directly denoting logical operations is typically late. In particular, developmental data directly pertaining to semantic acquisition of Hebrew conjunctions and disjunctions (coordination constructions) is delayed to age 5 and beyond. According to Paltiel-Gedalyovich 2003 children can use coordination structures correctly earlier in life, but due to processing demands they fail to do so consistently. This is precisely the time course reported for the acquisition of adjectives, in accordance with the thesis that adjectival dimensions are processed as rules (conjunctions / disjunctions) and consistent use of rules maturates at age 5.

Going beyond acquisition, the noun-adjective distinction turns out to have important implications with regard to inference. For example, subjects tend to draw conditional inferences from nouns to adjectives (as in *if a circle, then red*, more easily than vice versa (as in *if red, then a circle*; Fugardy et al. 2009). Naturally, a averaged value in multiple dimensions can hardly be inferred from a value on a single dimension, but a value on single dimension may be inferred from a high/low weighted sum/product.

Furthermore the noun adjective distinction may have neural correlates, and studies of the latter have bearing on the RS hypothesis.

One type of data comes from studies of aphasia and dyslexia. Numerous reports exist of patients whose speech, reading or writing performance is differentially affected for nouns, verbs, and adjectives. The typical result reported is for better performance with nouns relative to verbs and adjectives. For example, Irigaray (1973) found reduced adjective rates in probable dementia of Alzheimer type participants by comparison with normal controls; In another study of aphasics, McNeil, Doyle, Spencer, Goda, Flores and Small (1997), asked an aphasic subject to produce a

synonym or antonym for a verbally presented word; a cueing treatment applied upon failure to produce a correct response. The subject reached criterion for production of noun synonyms, but failed to reach criterion for adjective synonyms. Another report by Coltheart, Patterson and Marshall (1980) concerns a dyslectic patient misreading adjectives as related nouns (e.g., *wise* as *wisdom*; *strange* as *stranger*, etc.) or vice versa (e.g., *truth* as *true*). Potentially relevant to us is the fact that the performance of this patient with concrete nouns is reported to be better than with adjectives, verbs and also abstract nouns.

Concerning localizations, recall that rule-based ('Boolean') processing tasks demand the use of working memory and executive functions, thus recruits primarily frontal-striatal regions. Conversely, similarity tasks involve implicit procedural learning, thus typically recruits regions which do not project directly to the prefrontal cortex (Ashby and Maddox 2005: 163-165). Returning to nouns and adjectives, studies of naming tasks with object categories which are typically nominal, such as animals and artifacts, accord with the RS hypothesis; temporal, rather than frontal activation is typical of the tested nominal categories (Bookheimer 2002; Martin 2003: 66-67, Cappelletti et al. 2008: 717; Shapiro et al. 2006). Shapiro et al. (2006) report event-related responses specific to noun production, which were replicated with real and pseudo nouns, abstract and concrete nouns, and regular and irregular morphology. They hypothesize that nominal temporal region activations support the retrieval of lexical information including dimensional clusters (whereas left inferior frontal activations in verbs contribute to syntactic unification and motor action processing).

Furthermore, interestingly, studies of nominal categories report left frontal activation (BA 45,47) in similarity and abstract/concrete judgments (Martin 2003: 66-67). This result might rest on the fact that these tasks involve adjectival processing; selective attention to some of a item's features (e.g., its value along the *abstract* vs. *concrete* dimension) is demanding.

This indirect support for the RS hypothesis suggests that future research may profit from studies directly comparing grammatical processing of adjectival vs. nominal constructions, starting with the neuro-cognitive model proposed by Ashby and Maddox (2005) for rules and similarity-based categories. Such studies are necessary to support or refute the hypothesis of a more significant involvement of frontal regions in the activation of adjectives than nouns, as detailed above (Ashby and Maddox 2005: 163-165).

To summarize:

18) a. Rules are easy to reason about explicitly. Conversely, similarity-based criteria are hardly accessible through introspection.
    b. Explicit, rule-based processing is more demanding, requiring working memory and executive functions, which recruit frontal-striatal circuits. Conversely, Similarity involves implicit procedural processing.
    c. Children younger than 3 years old consistently perform similarity-based processing tasks. Conversely, due to the late maturation of the prefrontal cortex, by age 3 children tend to still have difficulties in consistently using rules, which only get resolved at age 5 (Frye et al. 1995; Zelazo et al. 1996, 2004).

In accordance, the acquisition of adjectives lags behind that of nouns and is characterized differently; the number and sort of dimensions affect acquisition rate, where fewer dimensions and logical dimension integration result in learning

difficulties. Thus, psycholinguistic data is consistent with the RS hypothesis. While similarity-based processing is the natural option for nouns, rule-based processing (which is difficult to young children) appears to be a dominant option for adjectives. An advantage of the RS hypothesis is that it provides clear neural and developmental predictions to start with, as well as well-studied paradigms to study them (see discussion and references in Ashby and Maddox 2005).

3 The domain of application of the RS hypothesis within and across languages

The third and final stage of the present study addresses non-paradigmatic subclasses of nouns and adjectives. Such subclasses challenge theories of word classes. Let us see how the RS hypothesis fairs with respect to them.

The non paradigmatic cases are systematic. They include nouns that are morphologically and/or semantically based on adjectives, or adjectives that are morphologically and/or semantically based on nouns. Moreover, the cases addressed below exhaust the range of examples I am aware of, which appear problematic for the generalizations discussed in this paper. They subsume the examples noted to me by reviewers, conference participants, and colleagues, for which I am most grateful.

3.1 Nouns which are directly related to an adjective

Two types of nouns systematically license *with respect to* (wrt) modification, namely animate-evaluative nouns like *an idiot* and property-denoting nouns, which are often called nominalizations, like *happiness, height, health, success, agreement, similarity*, and *difference*. These nouns are similar to adjectives in other respects as well, such as agreement and argument structure, as explained shortly.

Nominalizations form exceptions to almost any generalization about nouns, whether syntactic or semantic. For example, usually verbs or adjectives denote event types (categories of events or states), while nouns denote non-eventual entity-types, and, usually, verbs or adjectives have an elaborate argument structure, while nouns do not (Landman 2000). Nominalizations are atypical nouns in that they often denote event types (*success, disaster, agreement*) and they have an elaborate argument structure. For instance, while normal nouns cannot take *for* arguments, as in *# Tweety is a bird for a water-bird*, these exceptional nouns do, as in *the conference was a success for a student conference*. [14]

Similarly, nouns with an animate-evaluative component, like *idiot* and *genius* are atypical nouns. For example, in languages like Hebrew, the morphological form of verbs and adjectives (19a), but usually not of nouns (19b), agrees with the subject in gender. In addition, the copula can be omitted when predicate position hosts an adjective (19a), but usually not when it hosts a noun (19b). The animate-evaluative nouns behave like adjectives (19c).

19) a. *Dan (hu) yarok* [Dan is green$_{MASC}$]; *Beth (hi) yeruka* [Beth is green$_{FEM}$]
b. *Dan #(hu) cipor* [Dan is a bird]; *Beth #(hi) cipor* [Beth is a bird]
c. *Dan (hu) idiot* [Dan is an idiot$_{MASC}$]; *Beth (hi) idiotit* [Beth is an idiot$_{FEM}$]

Accordingly, these nouns license wrt-phrases, as in *Dan is an idiot wrt money / in every respect / except wrt money*, and *the conference was a success wrt the quality of the papers / in every respect / except for the papers.*

Concerning nominalizations, a line of investigation for the future is that the ability to combine with wrt- and for-phrases is due to the fact that adjectives that combine with a wrt-phrase can then be nominalized yielding a nominalization with a wrt-phrase. Importantly, in nominalizations the wrt-argument functions differently. It neither adds a categorization criterion as it usually does with adjectives, nor does it reduce the number of ordering criteria as it usually does.

For instance, compare *healthy wrt bp* with *health wrt bp*. Intuitively, an entity x falls under the adjective iff x possesses enough of the quality 'health wrt blood pressure' – x's blood pressure is close enough to the ideal. The wrt-phrase not only turns *healthy* into a one-dimensional adjective, but also adds a categorization criterion. In contrast, for x to fall under the corresponding noun, *health wrt bp*, x's blood pressure needs not be close to the ideal. Hence, the wrt-argument does not contribute to the noun the categorization criterion it contributes to the adjective; rather, x itself has to **be** that thing of which the adjective requires its instances to have enough, e.g. health wrt blood pressure. How do we decide whether something counts as *health*, some type of *malady*, *happiness*, *success* etc.? A rich set of symptomatic non-necessary dimensions characterizes disease types, and *success* can hardly be defined by a set of necessary conditions which are jointly sufficient. Thus, nominalizations are mapped to the noun category because they are not associated with a single adjectival respect; rather, they are associated with a rich set of typicality dimensions. To fall under the denotation of a nominalization, an entity has to resemble closely its prototype.

In line with this observation, extensive work by Moltmann (e.g., 2010) suggests that all nominalizations are characterized by multiple similarity dimensions, as the RS hypothesis predicts, given their syntactic category. Moltmann shows that we never judge as felicitous statements of the form "NP_1's Nominalization is NP_2's Nominalization", as in #*John's height is Bill's height*; rather, we say that *John's height is the same as Bill's height*. The use of *same* helps picking out the one property along which the two heights are identical, from all other properties along which they may be different. Hence, nominalizations, like other nouns denote entities, and are associated with multiple non-definitional dimensions of these entities. Hence, the classification of nominalizations as nouns is consistent with the RS hypothesis.

Similarly, concerning animate-evaluative nouns, a line of investigation for the future is, again, that they come to be when adjectives combined with a wrt-phrase are nominalized. Many evaluative nouns have an adjectival homonym. As adjectives they can be modified by *more* and *wrt*, and they can denote different ontological types of entities – human beings, cities, dishes. etc., as in *Dan is more Italian than Sam wrt their cooking* and *Florence is more Italian than Torino wrt food and weather*. The nominal homonym is more restricted. It can only denote human beings, as in *Dan is an Italian* vs. *# Florence is an Italian.* Hence, once a noun, this predicate is associated with a richer set of dimensions, as the RS hypothesis predicts. The noun phrase *an Italian* is interpreted more on a par with the modified phrase *an Italian man*, suggesting the presence of an adjective, *Italian*, in the morpho-syntactic structure, or derivation, as well as in the dimension set.

Some animate-evaluative nouns in English do not have an adjectival homonym, but are semantically associated with a scalar, adjectival dimension. For example, *genius* and *idiot* are associated with scalar dimensions of *intelligent*. Similarly, empirical research shows that the processing of categories such as those denoted by *boy, girl* and *child* is affected by gradable similarity dimensions such as *childish look and behavior*. In languages such as Spanish and Hebrew these nouns combine with

adjective-selecting degree morphemes like *very* and *more*, yielding an interpretation which is precisely identical to their adjectival dimension, e.g., *childish* and *intelligent*. However, these combinations are not as smooth as combinations with noun-selecting modifiers such as *really*. I propose that in those languages they have adjectival homonyms after all, although highly infrequent ones. This stipulation is a low price to pay, considering the robustness of the licensing restrictions on adjective-selecting degree morphemes like *very* and within-predicate *more*, which are incompatible with all other nouns, including nominalizations. Examples like *Dan is more Italian than Sam* become odd when an article is added, as in *#Dan is more an Italian than Sam*, unless the particle *of* is added. In addition, my searches of the whole internet, I have found but few examples of nouns combined with *more* independently of the morpheme *of*.

20) a. That's how much more a success Torino was, compared to Athens
    b. I'm always a boy; but I'm more a boy when I perform
    c. I'm more a boy than everyone in your team

In all the other cases, comparisons were between degrees in two different predicates.

    d. Probably this is more an Italian tradition than a British one
    e. To Italians he is almost more an Italian than an English poet
    f. These young Japanese Americans prove their patriotism through unquestioning obedience to authority, ironically a trait more Japanese than American
    g. Columbus was more a "success" for having landed in the Bahamas than in Bombay
    h. He's much more a boy from Long Island than a boy from Brooklyn
    i. The hero seemed more a boy than a man.

Hence, these nouns too are sharply distinguished from adjectives, for which plenty of examples can be found of within-predicate comparisons:

    j. The southern part of the region is far more Italian than Alto Adige
    k. More charming than Boston, more romantic than Vegas and more Italian than Naples, Providence is an undiscovered gem of a city with no traffic.

Recall also that the RS hypothesis readily provides a systematic way by which nouns can be derived from adjectives, and vice versa. According to this thesis, both nouns and adjectives are associated with dimensions, and they are distinguished minimally by the ways the dimensions are integrated. Hence, despite appearances to the contrary, granting the existence of infrequent adjectival homonyms for nouns like *genius* and *idiot*, the data is consistent with the RS hypothesis.

3.2 Adjectives derived from nouns

Some languages have noun-like adjectives, whereas others have verb-like ones, a distinction expressed by the type of degree constructions available in a language.[15] Furthermore, languages with a well established adjective set exhibit variance with regard to the syntactic classification of predicates as adjectives and nouns. This section discusses two cases – French relational adjectives, which are morphologically

derived from nouns, and the Spanish Det-A construction, in which a determiner and an adjective occur with no overt noun. Last, but not least, the cross-linguistic rarity of adjectives is addressed.

3.2.1 *Relational adjectives*

Relational adjectives form yet another non-paradigmatic class (Rodriguez Pedreira 2000, Schuwer 2005, McNally and Boleda 2004).[16] Focusing on French examples like those in (21), we see that many of these adjectives classify as nouns in English, where they can function as modifiers in noun-noun compounds.

21) a. Une carte routière (a road map);
b. Un régime présidentiel (a presidential regime); un voyage présidentiel (a visit of the President to …)
c. Une lampe halogène (halogen lamp)
d. La piscine municipale (the public swimming pool)
e. Le voyage alsacien du ministre (the minister's trip to Alsace)
f. Le cerveau humain (the human brain)
g. Le pays natal (native country)
h. Le lait maternel (mother's milk)
i. L'acné juvénile (teenage acne)
j. une spécialité régionale (a regional speciality)
k. tronc cérébral (brain stem)
l. une décision gouvernementale, un compte bancaire, un directeur administratif, la recherche scientifique, une réunion ministérielle….

This group of adjectives has the following characteristics in French. Some apply in English, too. First, a relational adjective comes *after* the noun it modifies (e.g., * *Une routière carte*), whereas non-relational adjectives can often be placed before or after the noun. Second, these adjectives are not normally used predicatively (22a,b), although for available predicative uses see McNally and Boleda's (2004) corpus study, and the contrastive construction which facilitates predicative use in (22c).

22) a.??Cette carte est routière (This map is a road map)
b. ?This regime is presidential (This is a presidential regime)
c. Cette critique n'est pas musicale mais littéraire (this critique is not a musical, but rather, a literary one)

Third, a relational adjective is not normally modifiable by a degree adverb (e.g., **Une carte très routière*), which is also true in English (**Her very/extremely/relatively native country;* ??*A slightly presidential regime*). Thus, like nouns, these adjectives are not morphologically gradable (cf. Section 1.2).

The above restrictions, which pertain to syntactic position and degree morphology, are problematic to hypotheses which base the adjective-noun distinctions on syntactic position and degree morphology, respectively. However, they are consistent with the RS hypothesis, which does not rest directly on either. Non gradable adjectives may count as adjectives, providing that their binary dimensions are processed as rule based (e.g. *prime; even; native*); the dimensions may be highly context dependent, but should, by default, function as categorization criteria (e.g., *presidential*).

What is more, the more academic, formal and technical a register, the more frequent the occurrence of relational adjectives instead of nouns used attributively. There are notably many relational adjectives in academic English as well; for instance, we find *renal failure* next to *kidney failure*, *dental decay* next to *tooth decay*, *pulmonary cancer* next to *lung cancer*, etc. Naturally, academic, medical, and technical endeavours involve an attempt to define words explicitly and precisely, namely through categorization criteria rather than mere similarity features.

A hypothesis whereby nouns and adjectives are all alike semantically fails to explain this tendency to coin adjectives where natural language speakers not doing science or technology coin nouns. However, the RS hypothesis predicts precisely that. Contextual demand for high standards of precision, as in scientific and technical contexts, biases towards rule-based processing, whereas contextual demand for fast processing and low precision level (approximation) biases towards similarity-based processing. Future research should address the interaction between contextual biases and category-biases for rules vs. Similarity (see also Sassoon 2011b).

3.2.2 Adjectives without nouns

In Spanish, phrases consisting of a determiner followed by an adjective can occur with no overt noun. Unlike English, cf. (23-25a), or even French, cf. (23-25b), in Spanish this process is pervasive, and not restricted lexically (23-25c).

23)	a. the poor	b. les pauvres	c. los pobres
24)	a. *the smooth	b. *le lisse	c. el suave
25)	a. *the careful	b. *les prudents	c. los cuidadosos

This construction is sometimes attested in English corpora, such as the BNC, as in *While **the powerful** seem to get away with serious crimes, **the powerless** commit less serious offences and get prison.* However, as a general case, English speakers differ from Spanish speakers; they normally regard such examples as marked.

Det-A constructions typically refer to nominal object categories. The ‘dropped’ noun in these expressions can be human or not, singular or plural, specific or generic. Its interpretation is often roughly equivalent to the English *type* or *one*. The determiner and adjective retain gender and number concord with the elided noun.

With regard to acquisition, a study of CHILDES (Waxman, Senghas, and Benveniste 1997) reports that children of age 2;8 and 3;5 produce the Spanish det-A construction. Hence, Spanish children learn that adjectives, like nouns, may be used to reference object categories. Three to seven year-old Spanish speakers extend novel words presented as either count nouns or adjectives on the basis of object-category communality (Waxman, Senghas, and Benveniste 1997). Although this was less pronounced with novel adjectives than nouns, this inclination appeared to become stronger with age and increasing language experience. Also, 21–23 months old infants acquiring English successfully extend adjectives but not nouns on the basis of property-based commonalities such as color and texture, while infants acquiring Spanish do so later, between 23–29 months, when they continue to extend also based on object categories. Thus, children acquiring different languages appear to acquire slightly different tacit expectations regarding the range of application associated with adjectives, depending on the availability of adjectival constructions with an elided noun. For a typological map of languages allowing adjectives without nouns see Gil 2005a.

The RS hypothesis predicts nominal categorization to be based on multiple similarity dimensions, and categorization under adjectives to be based on single categorization criteria or conjunctions or disjunctions of several such criteria. To account for the apparent resemblance of adjectives to nouns in Spanish-like languages, the RS hypothesis predicts that it is either the case that all of the adjectives in these languages have nominal homonyms (which is intuitively the case in, for example, Hebrew), or alternatively, that the Det A construction hosts an implicit noun or pronoun (as the authors above assume), or both. Either possibility explains the difficulty of acquiring adjectival interpretations in these languages. If both options are discovered to be wrong, this will speak against the RS hypothesis. A third alternative, whereby argument position triggers similarity-based processing, while modifier position triggers rule based processing, does not explain differences between adjectives and nouns in, for example, predicate position.

3.2.3 *Languages with a closed set of adjectives*

While the noun category is relatively stable in the languages of the world, languages vary widely with regard to the existence and size of the set of adjectives. Consequently, the universality of the category Adjective has been questioned (Dixon 1982). Recently, Baker (2003) argued extensively against the view that some languages do not have adjectives at all, and hence their grammar does not employ such a syntactic category. All languages seem to have at least a small group of adjectives. For example, Igbo has a total of 8 adjectives, Hausa has 12, and Bantu languages have adjective classes ranging from ten to fifty words.

The prevalence of nouns in comparison with adjectives cross linguistically aligns with their earlier and faster acquisition (cf. section 2 above). On the RS proposal, this prevalence follows from a differential processing demand. The two categories can do the same sort of a thing – designate object categories. However, nouns do so by means of multiple criteria combined through similarity functions, while adjectives do so by means of single criteria or Boolean compounds of such criteria. Accordingly, acquisition of adjectives is more demanding than nouns. Thus, predictably, across languages, lexicons will include more of the less demanding option (nouns) than of the more demanding one (adjectives).

Moreover, the reasoning behind the RS hypothesis applies beyond the noun-adjective distinction, to any construction in any language which exhibits a rule based structure, namely the features discussed in section 1 – categorization criteria or conjunctions and disjunctions of such criteria, morphological gradability, i.e. compatibility with degree morphemes like *very, too, enough* and so on, and accessible dimensions, i.e. compatibility with prepositions which mark dimension arguments, like wrt-phrases, quantification over dimensions and dimensional exception phrases.

Considering, for illustration, English verbs, we see that they divide into two different groups. Some verbs resemble nouns in not being gradable, e.g., *kissed more* can only mean ‘was involved in more kissing events’ or ‘in a temporally longer kissing event’; it cannot mean ‘exemplifies better the property *kissed*’. Thus, an utterance of the sentence *Bill kissed Mary more than he kissed Sue*, if acceptable at all, is unlikely to convey attribution of a higher degree of ‘kissing-ness’ to one event than to another. Consideration of examples with other degree morphemes like *so, too,* and *as,* yields a generalization whereby morphological gradability with verbs like *kiss*, relates to measures of time duration or number of kissing events, not to degree.

Other verbs do exhibit gradability of the sort prevalent with adjectives, including for example measuring verbs such as *weighs, costs lasts,* and *resembles,* as well as stative psychological verbs such as *interests, fears, frightens, believes* and *loves*; e.g. an utterance of the sentence *Dan loves Mary more than he loves Sue* conveys attribution of a higher degree of love (rather than longer temporal duration), and *Dan resembles Mary more than Sue* conveys attribution of a higher degree of resemblance to Mary. In accordance, multidimensional gradable verbs allow for *with respect to* phrases, as in *Dan resembles Mary wrt the color of their skin and the form of their lips.* Verbs with multiple rule-based dimensions like *resembles* classify as multidimensional. Degree achievements like *to weigh* and psychological verbs like *to love* classify as one-dimensional, and *to kiss* – as similarity-based.

Hence, the rule vs. similarity distinction cross cuts the category of verbs. In accordance, some languages do not have many adjectives, and rather employ stative verbs, which resemble the English degree achievements, to grade and evaluate entities (e.g., Vanuato; see also Baker 2003: 249-263). This is consistent with the RS proposal, as nothing prevents a language from extensively using the possibilities allowed by the verbal category.[17] Other languages, like Chichewa, use nominalizations, as in the English *has height, more height, much height,* etc. (Baker 2003: 245-248). This is once again consistent with the RS proposal, as nothing prevents a language from extensively using the possibility of counting or measuring amounts of stuff denoted by a noun, as in the English *you brought us more sand than water* or *more apples than pears* (cf. the discussion in section 1.4).

4 Conclusions

The RS hypothesis provides a new and fruitful direction of research into the puzzling topic of word classes. It provides a coherent explanation of a wide set of facts pertaining to the structure of nominal and adjectival categories, namely, psychological and linguistic facts pertaining to the type of categorization criterion and degree function of nouns and adjectives, to the brain mapping and acquisition of these syntactic categories, and to a number of non paradigmatic sub-classes. This suggests that future research will profit from a more direct study of the status of the RS hypothesis.

***Acknowledgments**: The ideas developed in this paper originate from work on my diss. (Sassoon 2007: Chap. 7), which was made possible by the Orgler Scholarship, Tel Aviv University (2004-2007). Part of the research for this paper was carried out in the project `On vagueness – and how to be precise enough', funded by the Netherlands Organization for Scientific Research (NWO 360-20-201). I warmly thank my PhD supervisor Dr. Nirit Kadmon, my teacher Prof. Fred Landman and the audience in IATL25 (Ben Gurion University of the Negev), UICM3 (Brussels) and WC2010 (Roma-Tre) for their helpful comments. Special thanks to Bart Geurts and Adar Weidman for insightful remarks on a manuscript that has eventually evolved into section 1, and to Philippe De Brabanter and two anonymous reviewers for their helpful comments. Any mistakes are solely mine.

Notes

1. For a different view, see Gil 2000 and Haspelmath et al. 2005.
2. Some readers may favor an approach whereby nouns denote natural kinds, and are rigid designators; this is orthogonal to the point being made here for the following reasons. First, on this view, nouns associate with entity sets through a kind-realization relation, which is sensitive to context. Second, there is hardly an agreement about what constitutes a 'kind', and it is far from clear whether all nouns and no adjectives (nor noun-adjective combinations) should be associated with a kind. Hence, to provide a general motivation for the noun-adjective distinction in terms of an association with a

kind or its absence thereof, a specific notion of a 'kind' has to be provided that is general enough to cover all nouns, but not too general so as not to cover all other concepts.

Notice also that the paper does not assume that the reference of nouns proceeds by the satisfaction of some condition or predicate, rather than relationally, through some causal relation between the noun or the speaker and the reference. Nouns and adjectives alike are assumed to be associated with a referent in the world, as well as with a classification condition. The paper does not take a stand as to which of these two is more constitutive of meaning.

3. For a review of this work see Mervis and Rosch (1981). For a review of earlier studies which form the basis for this work see Lakoff (1987; chapter 2). Reviews of more recent developments and theoretical approaches are found in, for instance, Hampton (1997) and Murphy's (2002) seminal "big book of concepts".
4. The origins of the exemplar theory are found in Medin and Schaffer's (1978) context model. Exemplar models include Hintzman's (1986) Minerva model; Nosofsky's (1988) generalized context model; Kruschke's (1992) ALCOVE model; Estes's (1994) array model, etc.
5. Exemplar models extend the number of representations which we encode in memory. For instance, it is assumed that, for a concept like 'bird', we encode in memory separate dimension sets for 'robin', 'duck', 'chicken', etc. It is predicted that if an item is highly similar to at least one known exemplar, it is highly similar to the concept. In addition, concepts like 'bird' are assumed to belong to a set of "contrast-concepts" (that includes concepts like 'mammals', 'reptiles', 'insects', etc.) Items are assumed to be classified in the contrast concept to which they resemble most (Ashby and Maddox 1993).
6. Unfortunately, a proper review of semantic theories of gradability and comparison is beyond the space limits of this paper. Detailed reviews are found in the third volume of the *Journal of Semantics* (especially, von Stechow 1984), Klein (1991) and Kennedy (1999).
7. Exceptional formal semantic theories assigning nouns gradable interpretation are Kamp and Partee 1995 and more recently Morzycki 2009.
8. Notice also that nominal examples such as #*is a bird with respect to…* seem better off if *bird* is interpreted metaphorically, as conveying the adjectival meaning *bird-like* ("similar to a bird wrt…").
9. We can utter statements like *Non-Japanese who love Japan become more Japanese than the Japanese* also because the noun and adjective denotations need not be identical in virtue of the indeterminacy in the dimension set of the adjective and to some extent also the noun. In the given example, the noun denotes the set of Japanese by nationality or birth, while the adjective is interpreted wrt behavior (the behavior of the non-Japanese by birth is more of a stereotypical Japanese behavior than the behavior of the Japanese by birth). That is, statements like *Dan is Japanese* need not relate to a completely identical dimension set as ones like *Dan is a Japanese*, if, for instance, *Japanese* in the former statement is interpreted *wrt stereotypical Japanese behavior*. This holds true of nationality adjectives and nouns generally, beyond a particular morphological form (cf. *(an)American; (an) Italian; (an) Israeli*, etc.), although in some cases, the additional +human restriction is made overt in the nominal form (cf. *French* vs. *Frenchmen*).
10. Notice also that nouns can function as adjectives merely by virtue of occurrence in modifier position, which is typically adjectival (as in, for instance., *an elephant turtle*; crucially, e.g., *a turtle which is an elephant* should be something that is both a turtle and an elephant, while *an elephant turtle* is simply a turtle with some dominant elephant feature.
11. Color terms, specifically, are acquired later than many other property terms; according to Bornstein 1985 this results from dependency on the maturation and integration of cortical neurological structures specific to color naming.
12. Similar observations about the interpretation of typicality modifiers were given by McCready and Ogata (2007). However, the current proposal is novel in that it derives these interpretations from a basic interpretation rule for adjectives in general.
13. In principle, another interpretation for "#Dan is not healthy (in every respect) except bp" could have been available, whereby negation outscopes the (implicit) universal quantifier, as in "it is not the case that Dan is healthy in all respects except bp". However, this reading is not available, except perhaps with a very special intonation, meaning that the distribution of exception phrases is restricted to 'positive' (upward entailing) contexts (von Fintel, 1994). Thus, indeed, no exception phrases are expected to occur with negated conjunctive adjectives, especially in written corpora.
14. Notice that the denotation of some nominalizations (say, *height*) may be fixed, but the denotations of many others (*success*; *health*) may be highly context dependent.
15. See Stassen 2005b for a typological map of comparison structures in the world languages; for formal characteristics distinguishing adjectives from verbs in the languages of the world see Stassen

2005a; see Gil 2005a for a discussion of how languages vary with respect to grammatical encoding of different types of attribution (genitive, adjective and relative clause constructions).

16. I am deeply grateful to Philippe De Brabanter for bringing these adjectives to my attention and teaching me about their intriguing characteristics.
17. It may be interesting to test whether it is possible to extend the RS hypothesis to distinguish between verbs and adverbs, since some distinctions between these two word classes parallel the distinctions between nouns and adjectives.

References

Ashby, Gregory, F., and Maddox, W. Todd. 1993. Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology* 37: 372-400.

Ashby, Gregory, F., and Maddox, W. Todd. 2005. Human Category Learning, *Annual Reveiew of Psychology* 56:149-78.

Baker, Mark C. 2003. *Lexical categories, Verbs, Nouns and Adjectives*. Cambridge Studies in Linguistics, Cambridge University Press.

Bartsch, Renate. 1986. Context-dependent interpretations of lexical items. In: Jerome Groenendijk, D. de Jongh and Martin Stokhof (eds.) *Foundations of Pragmatics and Lexical Semantics, GRASS 7*, Foris, Dordrecht.

Berman, Ruth. 1988. Word class distinctions in developing grammars. Categories and processes in language acquisition, Hilsdale NJ., Erbbaum.

Bookheimer, Susan, (2002) Functional MRI Of Language: New Approaches To Understanding The Cortical Organization Of Semantic Processing. In: Annual Review Of Neuroscience 25, 88-151.

Booth, A.E. and Waxman, S.R. 2009. A horse of a different color: Specifying with precision infants' mappings of novel nouns and adjectives, *Child Development* 80(1): 15-22.

Bornstein, M. H. 1985. On the development of color naming in young children: Data and theory. *Brain and Language* **26**, pp. 72–93.

Cappelletti, Marinella, Felipe Fregni, Kevin Shapiro, Alvaro Pascual-Leone, and Alfonso Caramazza. 2008. Processing Nouns and Verbs in the Left Frontal Cortex: A Transcranial Magnetic Stimulation Study. *Journal of Cognitive Neuroscience*. 20(4): 707-720.

Carlson, Gregory. 1977. *Reference to Kinds in English.* Doctoral Dissertation, University Of Massachusetts, Amherst (Also Published In 1980, New York, Garland).

Chater, N., & Hahn, U. (1997). Representational distortion, similarity, and the universal law of generalization. *Proceedings of the Interdisciplinary Workshop on Similarity and Categorization*, SimCat97: 31–36.

Edinburgh: Department of Artificial Intelligence, University of Edinburgh.

Coltheart, M., Patterson, K., Marshall, J., (eds.) 1980. *Deep Dyslexia.* London: Routledge & Kegan Paul.

Damon, William, Lerner Richard M., Kuhn Deanna and Siegler Robert S. (eds.) 2006. *Handbook of child psychology*, volume 2 *Cognition, Perception and Language*. 6th edition, John Wiley & Sons.

Dayal, Veneeta. 2004. Number Marking and (In)definiteness in Kind Terms. *Linguistics and Philosophy* 27(4): 393 – 450.

Dixon R. M. W. 1982. *Where have all the adjectives gone?* Mouton, Berlin.

Estes, William, K. 1994. *Classification and cognition*. Oxford. Oxford Uni. Press.

Evans, N. 2000. Word classes in the world's languages. In G. Booij, C. Lehmann, & J. Mugdan (eds.), *Morphology: A handbook on the inflection and word formation*. Berlin.

Fodor, Jerry A., Garret M., Walker E. and Parkes C. 1980, Against Definitions. *Cognition* 8, 263-367.

Frye, D., Zelazo, P. D. & Palfai, T. 1995. Theory of mind and rule-based reasoning. *Cognitive Development 10*, 483-527.

Fugardy, Andrew J. B., Niki Pfeifer, Bastian Mayerhofer, and Gernot D. Kleiter. 2009. How people interpret an uncertain *If*, MS, University of Salzburg, Austria.

von Fintel, Kai. 1993. Exceptive constructions. *Natural Language Semantics* 1.2:123-148.

von Fintel, Kai. 1994. *Restrictions on Quantifier Domains*. Ph.D. Dissertation, Uni. of Massachusetts, Amherst.

Gassar, Michel and Linda Smith. 1998. Language and *cognitive processes* 13(2/3): 269-306.

Gibson, Eleanor J. 1969. *Principles of conceptual learning and development*. New York. Appleton Century Crofts.

Gil, David. 2000. "Syntactic Categories, Cross-Linguistic Variation and Universal Grammar". In P. M. Vogel and B. Comrie eds., *Approaches to the Typology of Word Classes, Empirical Approaches to Language Typology*, Mouton, Berlin and New York, 173-216.

Gil, David. 2005a. Genitives, Adjectives and Relative Clauses, *The World Atlas of Language Structures online*, Chapter 60: Oxford University Press.
Gil, David. 2005b. Adjectives without Nouns. *The World Atlas of Language Structures online*, Chapter 61: Oxford University Press.
Gvozdev, A. N. 1961. Voprosy izučenija detskoj reči. Moscow: Akademija Pedagogiceskix Nauk, RFSFR.
Hampton, James. 1979. Polymorphous concepts in semantic memory, *Journal of Verbal Learning and Verbal Behavior* 18: 441-461.
Hampton, James. 1995. Testing the prototype theory of concepts. *Journal of Memory and Language*. 34: 686-708.
Hampton, J.A., 1995. Testing prototype theory of concepts. *Journal of Memory and Language* 34, 686–708.
Hampton, James. 1997., Conceptual Combination. In Koen Lamberts and David Shanks (eds.), *Knowledge, Concepts and Concepts*. The MIT Press. Cambridge, MA: 135-162.
Hampton, James. 1998. Similarity based categorization and fuzziness of natural concepts. *Cognition* 65: 137-165.
Haspelmath, Martin, Matthew Dryer, David Gil and Bernard Comrie (eds.) 2005. *The World Atlas of Language Structures*, Oxford University Press, Oxford.
Hillis, A. E. and Rapp, B. (eds.) 2001. The organisation of the lexical system. *The handbook of cognitive neuropsychology*: 185-210. Psychology Press, Philadelphia, PA.
Hintzman, Douglas, L. 1986. 'Schema abstraction' in a multiple-trace memory model. *Psychological Review* 93: 411-28.
Irigaray, L. 1973. Le language des déments (The Hague: Mouton).
Kadmon, Nirit, and Fred Landman (1993). Any, *Linguistics and Philosophy* 16: 353-422.
Kratzer, Angelika (1986). Conditionals. In: A.M. Farley, P.T. Farley and K.E. McCulluough (eds.) *CLS 22 part 2: Papers from the parasession on pragmatics and grammatical theory*. CLS, University of Chicago.
Kamp, Hans. 1975. Two theories about Adjectives. In: Edward Keenan (ed.), *Formal Semantics for Natural Language.* Cambridge University Press: 123-155.
Kamp, Hans and Barbara Partee. 1995. 'Prototype theory and compositionality'. Cognition 57(2), 121–191.
Keil, Frank. 1989. *Concepts, kinds and cognitive developments*, The MIT Press. Cambridge, MA.
Kennedy, Christopher. 1999. *Projecting the adjective: the syntax and semantics of gradability and comparison*. Garland. NY. (1997 UCSC PhD thesis).
Klein, Ewan. 1980. A semantics for positive and comparative adjectives. *Linguistics and Philosophy* 4: 1–45.
Klein, Ewan. 1991. Comparatives. In: Arnim von Stechow and Dieter Wunderlich (eds.) *Semantik / Semantics, An International Handbook of Contemporary Research*, 673-691. Berlin, New York: de Gruyter.
Kellenbach, M. L., Wijers, A. A., Hovius, M., Mulder, J., & Mulder, G. (2002). Neural differentiation of lexicosyntactic categories or semantic features? event-related potential evidence for both. J of cognitive Neuroscience, 14(4), 561-577.
Kruschke, John. 1992. Alcove: An Exemplar Based Connectionist Model of Concept Learning. *Psychological Review* 99, 22-44.
Kripke, Saul A. (1972). Naming and necessity. In G. Harman and D. Davidson (Eds.), *Semantics of Natural Language*. Kluwer, Dordrecht.
Lakoff, George. 1987. *Women, Fire and Dangerous Things: What Concepts Reveal about the Mind*. Chicago University Press.
Landman, Fred. 1989. Groups II, *Linguistics and Philosophy* 12: 723-744.
Landman, Fred. 2000. *Events and Plurality*. Kluwer. Dordrecht: 143-164.
Lapata, Mirella and Frank Keller. 2005. Web-based Models for Natural Language Processing. ACM Transactions on Speech and Language Processing 2:1, 1-31.
Markman, E. M., & Jaswal, V. K. (2004). Acquiring and using a grammatical form class: Lessons from the proper-count distinction. In D. G. Hall & S. Waxman (Eds.), Weaving a lexicon (pp. 371-409). Cambridge, MA: MIT Press.
Martin, Randi. 2003. Language processing, Functional organization and Neuroanatomical basis, *Ann Rev Psy* 54: 55-89.
McCloskey, Michael, and Glucksberg, Sam. 1978. Natural concepts: Well defined or fuzzy sets? *Memory and cognition* 6: 462-472.

McCready, Eric, and Norry Ogata. 2007. Adjectives, stereotypicality, and comparison, *Natural Language Semantics* 15: 35-63.
McNally, Louise and Gemma Boleda. 2004. Relational adjectives as properties of kinds, In: O. Bonami & P. Cabredo Hofherr (eds.) *Empirical Issues in Formal Syntax and Semantics* 5, 179–196.
McNeil MR, Doyle PJ, Spencer K, Goda AJ, Flores DK, Small SL. 1997. A double-blind, placebo-controlled study of pharmacological and behavioural treatment of lexical-semantic deficits in aphasia. *Aphasiology* 11 (4,5): 385–400.
Medin Douglas, L. and Schaffer M. 1978. Context Theory of Classification Learning. *Psychological Review* 85: 207-238.
Mervis, C. B., and Rosch, E. 1981. Categorization of natural objects. In M. R. Rosenzweig & L. W. Porter (Eds.), *Annual Review of Psychology* 32.
Miceli, G. , Turriziani, P. , Caltagirone, C. , Capasso, R. , Tomaiuolo, F. and Caramazza, A. 2002. The neural correlates of grammatical gender: An fMRI investigation. *Journal of Cognitive Neuroscience* **14** , pp. 618-628.
Miozzo, Michele, Albert Costa, Mireia Herna´ndez and Brenda Rapp. 2010. Lexical processing in the bilingual brain: Evidence from grammatical/morphological deficits, *Aphasiology* 24(2), 262–287.
Moltmann, Friederike. 2010. Gradability within a trope based approach. Handout for the workshop *The representation of Degrees*, June 6-7, Leiden.
Morzycki, Maxin. 2009. 'Degree Modification of Gradable Nouns: Size Adjectives and Adnominal Degree Morphemes'. *Natural Language Semantics* 17(2), 175–203.
Murphy, Gregory. 2002. *The big book of concepts.* The MIT Press. Cambridge, MA.
Nosofsky, Robert, .M. 1988. Similarity, frequency, and concept representation, *Journal of Experimental Psychology: Learning, Memory and Cognition* 14: 54-65.
Partee, Barbara. 1989. Binding implicit variables in quantified contexts. In *Papers from the 25th Regional Meeting, Chicago Linguistic Society. Part 1. the general session*. Chicago.
Paltiel-Gedalyovich, Leah R. 2003. Towards an explanation of first language acquisition of Hebrew coordination, PhD Diss., Ben Gurion University.
Putnam, H. (1975). The meaning of 'meaning'. In Keith Gunerson (Ed.) *Minnesota Studies in the Philosophy of Science* ,Vol. VII, University of Minnesota Press, Minneapolis.
Polinsky Maria. 2005. Word class distinctions in an incomplete grammar. In Dorit Ravid and Hava Bat-Zeev Shyldkrot, eds. *Perspectives on language and language development*, 419-436. Dordrecht: Kluwer.
Pustejovsky, James. 1991. *The generative Lexicon*. Computational linguistics 17(4): 409- 441.
Rapp, B. and Goldrick, M. 2006. Speaking words: The contribution of cognitive neuropsychological research. *Cognitive Neuropsychology* 23: 39-73.
Ravid, D., Tribushinina, E., Korecky-Kröll, K., Xanthos, A., Kilani-Schoch, M., Laaha, S., Leibovitch-Cohen, I., Nir, B., Aksu Koç, A., Dressler, W.U. and Gillis, S. (2010). The first year of adjectives: a cross-linguistic study of the emergence of a category. Poster presented at *the Child Language Seminar*, London, 24-25 June 2010.
Rodriguez Pedreira, Nuria. 2000. adjectifs qualificatifs et adjectifs relationnels: étude sémantique et approche pragmatique, Doctoral thesis, Universidad de Santiago de Compostela.
Rosch, Eleanor. 1973. On the internal structure of perceptual and semantic concepts. In T. E. Moore (ed.), *Cognitive Development and the Acquisition of Language*, Academic Press. New York.
Rosch, Eleanor and Carolyn Mervis. 1975. Family Resemblance: Study In The Internal Structure Of Concepts. *Cognitive Psychology* 7: 573-605.
Ross, John, R. 1973. Nouniness. In: Osamu Fujimura (ed.), *Three dimensions of linguistic research*. Tec Company, Ltd. Tokyo: 137-257
Roth, Emilie, and Edward Shoben. 1983. The effect of context on the structure of concepts. *Cognitive Psychology* 15: 346-378.
Sapir, E. 1921. Language. New York: Hartcourt, Brace and World.
Sassoon, W. Galit. 2007. Vagueness, Typicality and Gradability, A comprehensive semantic analysis. PhD dissertation, Tel Aviv University.
Sassoon, W. Galit. 2010. Measurement theory in linguistics. *Synthese* 174(1) : 151-180.
Sassoon, W. Galit. 2011a. A typology of multidimensional adjectives. MS under review, ILLC, University of Amsterdam.
Sassoon, W. Galit. 2011b. The double nature of negative antonymy. To appear in: *The Proceedings of Sinn und Bedeutung 16*, Utrecht.
Schuwer, Martin. 2005. Système des adjectifs non prédicatifs en français et en anglais: constantes et variations, *Cahiers de Lexicologie* 86: 85-103.

Shapiro, Kevin A., Lauren R. Moo and Alfonso Caramazza. 2006. Cortical Signatures of Noun and Verb Production, *Proceedings of the National Academy of Sciences of the USA* 103(5): 1644-1649.
Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science* 237, 1317–1323.
Smith, David, J., and Minda, John Paul. 1998. Prototypes in the mist: The early epochs of concept learning, *Journal of Experimental Psychology: Learning, Memory and Cognition* 24: 1411-1436.
Smith, David, J., and Minda, John Paul. 2000. Thirty categorization results in search of a model, *Journal of Experimental Psychology: Learning, Memory and Cognition* 26: 3-27.
Solt, Stephanie. 2009. The semantics of adjectives of quantity. PhD dissertation, City University of New York.
Syrett, Kennedy and Lidz. 2009. Meaning and context in children's understanding of gradable adjectives. *Journal of semantics* 27: 1-35.
Stassen, Leon. 2005a. Predicative Adjectives. *The World Atlas of Language Structures online*, Chapter 118: Oxford University Press.
Stassen, Leon. 2005b. Comparative Constructions. *The World Atlas of Language Structures online*, Chapter 121: Oxford University Press.
von Stechow, Arnim. 1984. Comparing Semantic Theories of Comparison. *Journal of Semantics* 3:1-77.
Waxman, Sandra R., Ann Senghas and Susana Benveniste. 1997. Cross-Linguistic Examination of the Noun-Category Bias: Its Existence and Specificity in French- and Spanish-Speaking Preschool-Aged Children, *Cognitive Psychology* 32, 183–218.
Waxman, Sandra R. and Amy E. Booth. 2001. Seeing Pink Elephants: Fourteen-Month-Olds' Interpretations of Novel Nouns and Adjectives, *Cognitive Psychology* 43(3): 217-242.
Waxman, Sandra R. and Jeffrey Lidz. 2006. Early word learning. In D. Kuhn & R. Siegler (eds.), *Handbook of Child Psychology, 6th Edition*, Vol. 2: 299-335. Hoboken NJ: Wiley.
Waxman, S. R., & Kosowski, T. D. 1990. Nouns mark category relations: Toddlers' and preschoolers' word-leaming biases. Child Development, 61, 1461-1473.
Wittgenstein, Ludwig. 1968. Reprint: *Philosophical Investigations*, Translated by G.E.M. Anscombe, 3rd edition. Blackwell. Oxford. Original edition, 1953.
van Deemter, Kees. 2010. *Not Exactly: in Praise of Vagueness*. Oxford University Press.
Zelazo, P. D., Craik, F. I. M., and Booth, L. 2004. Executive function across the life span. *Acta Psychologica 115*: 167-184.

Appendix: *With respect to* in COCA 2010

As numerous examples from the corpus of contemporary American English (COCA 2010) illustrate, different uses of wrt-phrases exist, which are not dimensional-argument uses, i.e. not relating to a dimension of a word sense. These uses can be licensed with nominal concepts, perhaps even more easily than with adjectival ones. The appendix includes discussion of those.

As numerous examples from the corpus of contemporary American English (COCA 2010) illustrate, different uses of wrt-phrases exist, which are not dimensional-argument uses, i.e. not relating to a dimension of a word sense. These uses can be licensed with nominal concepts, perhaps even more easily than with adjectival ones.

One irrelevant use regards spatial or comparison relations. In this type of usage, wrt-phrases can be substituted with *relative to, compared to* or *in comparison with*.

(i) a. What is second wrt the third?
b. …is in a 30 degrees position wrt the horizontal axis
c. …high velocity with respect to earth

The third is not a dimension of interpretation of *second* (neither a categorization criterion nor a similarity dimension), and *the horizontal axis* is not a dimension of *30 degrees position* in the above sense. Nor is *earth* a dimension of *velocity*. Rather the arguments of these wrt-phrases – *the third, the horizontal axis* and *the earth* – are individuals or discourse entities similar to those denoted by the nominal concept modified by the wrt-phrase – *the second, 30 degrees position* and *high velocity*, respectively; the two arguments are said to be in some relation – a sequence relation, a spatial relation, a comparative relation of e.g. higher velocity, etc. In opposition, *blood pressure* does not stand in a spatial or comparative relation with *healthy*; rather, it functions as a categorization criteria of *healthy* in the sense that in order to count as healthy, one has to have normative blood pressure.

Another irrelevant use pertains to wrt-phrases marking a subject matter argument or a topic of discussion. In this usage, wrt-phrases can be substituted with *about, concerning, regarding* or *of*. Characteristic nouns with a subject matter argument include, for example, *problem, situation, question, hypothesis, position, opinion, proposal, rules, views, information, knowledge, roles, agreement, disagreement* etc.

(ii) a. During our tenure, there was the question wrt Macedonia of its diplomatic recognition…
b. Another is the situation wrt strategic weapon, which continues to be a matter of friction between us
c. That was the failure wrt 9/11
d. Where is the difficulty wrt the matter of the coherence of Smith's varying statements…

Strategic weapon is not a dimension of interpretation of *situations* – neither a categorization criteria nor a similarity dimension, and *9/11* is not a dimension of *failure* in this sense. At best, these are examples of situations/problems and failures, respectively. In opposition, *blood pressure* is not an example of a healthy individual, but rather a categorization criteria in the sense that in order to count as healthy, one has to have normative blood pressure.

Wrt-phrases are not common with other abstract nouns like *war* or relational nouns like *mother*, certainly not in a dimensional interpretation. In fact, considering COCA (2010), a balanced and annotated linguistic corpus, one finds three times as many uses of wrt-phrases with adjectives as with nouns in copula constructions, as in "is Adj. wrt" vs. "is a/an/□ noun wrt". The proportion of wrt-phrases among the "copula + noun" counts is .00006, while the proportion of wrt-phrases among the "copula + adjective" counts is .00018, which is three times as much. A more fine-grained analysis of the distribution and types of usage of wrt-phrases with adjectives vs. nouns awaits future research. Adding to the calculation definite noun counts of the form "is/are/was/… the noun wrt", does not affect the results significantly – the proportion of wrt-phrases among the "copula + noun" counts turns from .00006 to .00007, but importantly, all of the eight new hits are examples of irrelevant uses – in none does the wrt-phrase designate a noun dimension. This suggests that the actual ratio between the frequency of dimension uses in adjectives and their frequency in nouns is much bigger than 3 to 1.

The search for copula constructions is crucial, for in other hits of the form "noun wrt" or "adjective wrt" one cannot tell whether the wrt-phrase in fact modifies the predicate preceding it. Some counts of the form "noun wrt" may actually be uses in which the wrt-phrase modifies an adjective, as in, for instance, *Bill was notoriously healthier than his friends wrt to blood pressure*, or any other word, as in *significant differences among birds with respect to blood pressure*.