Formal Ontology in Information Systems
R. Ferrario and W. Kuhn (Eds.)

doi:10.3233/978-1-61499-660-6-167

Adjective Semantics in Open Knowledge Extraction

Aldo GANGEMI [a,b] Andrea Giovanni NUZZOLESE [b] Valentina PRESUTTI [b]
Diego REFORGIATO RECUPERO [b,c]

[a] *LIPN, Université Paris 13, Sorbonne Paris Cité CNRS UMR 7030, France*
[b] *Semantic Technology Lab, ISTC-CNR, Rome, Italy*
[c] *Department of Mathematics and Informatics, University of Cagliari, Italy*

Abstract. Adjectives carry a lot of discourse semantics, often being the core elements to understand the literal, emotional, cultural, and metaphoric meaning of a sentence. Their formal semantics has been investigated specially in lexical semantics and formal linguistics, with some contributions from formal ontology and the semantic web. However, no standard formal treatment of adjectives is available yet, which is capable to address the concerns of both theoretical and computational natural language understanding. In this paper we summarize the existing approaches, present some alternative solutions to approximate a rigorous but pragmatic representation, and describe an implementation of a lightweight adjective ontology as a core resource in FRED, a state-of-the-art open knowledge extraction tool.

Keywords. adjective semantics, knowledge extraction, formal ontology

1. Introduction

Adjectives carry a lot of discourse semantics, often being the core elements to understand the literal, emotional, cultural, and metaphoric meaning of a sentence. Their formal semantics has been investigated specially in lexical semantics and formal linguistics, with some contributions from formal ontology and the semantic web. However, no standard formal treatment of adjectives is available yet, which is capable to address the concerns of both theoretical and computational natural language understanding.

The main property that characterises an adjective is *sectivity*. More specifically, an adjective semantics have a different behaviour whether it is *intersective* or *subsective*. In the first case, for example *Canadian surgeon*, if we consider the individual referenced by the modified noun *surgeon*, and we hypothesise that she is also a pilot, we can entail that she is a "Canadian pilot". In the second case this entailment does not hold: for example if the phrase is *skilful surgeon*, we cannot entail that she is also a *skilful pilot*. Most adjectives are either intersective or subsective, nevertheless there are cases of ambiguity and other special types. An hypothesis that we analyse in this paper (cf. Section 2) is that sectivity may be explained by means of another property of adjectives named *framality*, which is the ability of an adjective to evoke a conceptual frame that is composed in a compatible way with other frames activated by related nouns and verbs in the same discourse. Compatibility refers to how (un)usual the composition of two or more frames is in a certain language/culture.

In order to support the proper interpretation of adjective semantics in open knowledge extraction, both representation patterns, and resources that distinguish them according to their semantic nature, are needed.

This paper provides two main novel contributions. The first is a set of patterns for representing adjective semantics (limited to *sective* and *framal* phenomena) in open knowledge extraction from discourse; the second is an ontology for about 7,000 English adjectives, automatically generated and evaluated, which extends WordNet and FrameNet adjective knowledge. The ontology and the representation patterns have been integrated into FRED [1] (a tool for open knowledge extraction), enabling a "safe" modelling of non-intersective adjectives.

The paper summarises the main problems related to adjective semantics (Sect. 2) and relevant related work (Sect. 3). Section 4 presents some solutions to an approximate but rigorous representation, while Section 5 describes the creation of the adjective ontology, and its integration as a core resource in FRED. Section 6 concludes the paper.

2. Main Problems in Adjective Semantics

As a part of his extensive work on modifiers, Morzycki [2] summarizes the central concerns in the semantics of adjectives, spanning through extensional vs. intensional adjectival modification, privativity, relationality, vagueness implicit in many gradable adjectives, etc. In order to keep the discussion simple and direct to the issues relevant to a formal ontology of adjectives, we reinterpret Morzycki's distinctions, augment them with new ones emerging from our empirical investigation, and single out three major aspects: (a) *sectivity*, (b) *framality*, and (c) *vagueness*. (a) and (c) are studied in depth in the literature, while (b) is a novel contribution of this paper. Here we extensively address (a) and (b).

2.1. Sectivity

Sectivity is a major criterion that distinguishes semantic phenomena involving adjectives. Sectivity impacts on the semantics of classes and individuals that are modified by adjectives. In practice, six types of phenomena can be singled out (cf. Sect. 4), and are introduced here via examples:

- INTERSECTIVITY: the adjective predicate that modifies a noun predicate can be independently predicated of the individual via entailment, e.g. $\forall(\mathtt{x})\mathtt{CanadianSurgeon(x)} \rightarrow \mathtt{Canadian(x)}$. In general, we can say that *Canadian surgeon* has an equivalent set-theoretic interpretation to the intersection of *Canadian* and *surgeon* interpretations. E.g. in description logic:

$$\mathtt{CanadianSurgeon} \equiv \mathtt{Canadian} \sqcap \mathtt{Surgeon} \tag{1}$$

  which is an exemplification of Heim and Kratzer's (cf. [3], [2]) *Predicate Modification Rule*, formalized here as the *Intersectivity Axiom Schema* (Axiom 2), where C denotes any phrase with adjectival modification, T the noun denoting the head/type of the multiword, M the modifying adjective, and $\cdot^I$ the interpretation of a constant in a domain $\triangle^I$:

$$C \equiv M \sqcap T \;,\; semantically: \; C^I = M^I \cap T^I \tag{2}$$

It seems that all proper adjectives are intersective.

- SUBSECTIVITY: the adjective predicate that modifies a noun predicate cannot be independently predicated of an individual via entailment, e.g. $\forall$`(x)SkilfulSurgeon(x)` $\nrightarrow$ `Skilful(x)`. Axiom 2 does not hold with subsective adjectives. An intuitive test for subsectivity can be performed via *type changing*. As Morzycki (cit.) exemplifies it, from *Floyd is a skilful surgeon* and *Floyd is an arsonist*, we cannot entail *Floyd is a skilful arsonist*, while the entailment holds if *skilful* is substituted by *Canadian*. The following is a clear example of how two subsective antonym adjectives can be jointly used, but no general entailment can be drawn: *he is a despicable husband, but an admirable doctor*.
- AMBIGUOUS SECTIVITY: the adjective predicate that modifies a noun predicate might be predicated or not of an individual via entailment depending on context, e.g. $\neg\forall$`(x)BeautifulDancer(x)` $\rightarrow$ `Beautiful(x)`. Adjectives with ambiguous sectivity cannot be assigned a definite interpretation without considering the discourse context, e.g. *beautiful* can be given an intersective interpretation in case there are other signs (for example gestures of appreciation or a previous comment), otherwise it is advisable to assign a subsective interpretation by default
- NON-EXTENSIONALITY: the adjective predicate can never be predicated of an individual, e.g. $\neg\exists$`(x)Alleged(x)` for $x \in \triangle^I$ with $\triangle^I \setminus V^I$, i.e. the set of individual constant interpretation in the domain of interpretation $\triangle^I$ is disjoint with the set of predicative constant interpretation in the vocabulary V^I.[1] For example, in the sentence *The alleged thief was arrested*, the individual predicated by the phrase *alleged thief* is not alleged, since we cannot substitute *thief* with e.g. *father*, and that individual is neither a thief, since no sentencing has been issued. Non-extensional adjectives are subsective, typically denote *modal* operators, and are often considered second order entities in the literature [4]
- SUBSECTIVE PRIVATIVITY: subsective adjectives that negate a core property of the noun predicate they modify, e.g. $\forall$`(x)(FakeGun(x)` $\nrightarrow$ `Fake(x)` $\wedge$ `FakeGun(x)` $\nrightarrow$ `Gun(x))`. However, the situation is not so neat, since not the whole interpretation of the noun predicate are negated by the privative, therefore we still need a (non-taxonomic) relation between e.g. *fake gun* and *gun*, and another with *fake*, as exemplified by the following description logic axiom:

$$\texttt{FakeGun} \equiv \exists\texttt{R.Gun} \sqcap \exists\texttt{S.Fake} \tag{3}$$

- INTERSECTIVE PRIVATIVITY: intersective adjectives that negate a core property of the noun predicate they modify, e.g. $\forall$`(x)(silver-tiger(x)` $\rightarrow$ `silver(x)` $\wedge$ `silver-tiger(x)` $\nrightarrow$ `tiger(x))`. The description logic formula for the example would be:

$$\texttt{SilverTiger} \equiv \texttt{Silver} \sqcap \exists\texttt{R.Tiger} \tag{4}$$

[1] Disjointness is required because non-extensional adjectives can be used to predicate over predicative constants.

For this usage of intersective adjectives, axiom 2 needs to be substituted with a non-taxonomic version, as in axiom 5.

$$C \equiv M \sqcap \exists R.T \ , \ semantically: \ C^I = M^I \cap \{x \mid for \ \langle x,y \rangle \in R^I, y \in T^I\} \tag{5}$$

As shown in the list, all six types of sectivity are variations on the applicability of the Intersectivity Axiom Schema. Sect. 4 will expand on those variations, suggesting alternative formal treatments for the different types.

2.2. Framality

Framality is defined here as the ability of an adjective to activate a conceptual frame (in Fillmore's sense [5]) during discourse interpretation, which may or may not compose with other frames activated by related nouns and verbs in the same discourse. Our hypothesis is that framality affects the way adjectives impact on the interpretation of discourse, and might be used as an explanation for sectivity phenomena.

First of all, some examples:

1. from `SkilfulSurgeon(x)` we cannot entail `Skilful(x)`, but from `ExtrovertedSurgeon(x)` we can entail `Extroverted(x)`, since *Extroverted* is intersective
2. from `BeautifulDancer(x)` we cannot safely entail `Beautiful(x)`, while from `BeautifulFox(x)` we can entail it
3. from `AuthoritativeScientist(x)` we cannot safely entail `Authoritative(x)`, but from `AuthoritativeWoman(x)` we can actually entail it

Why is this happening? What is in intersective adjectives like *extroverted* that enables this predicative independence? And why do certain adjectives become intersective with certain types (*fox* instead of *dancer*, *woman* instead of *scientist*)?

Framality is our proposal for a cognitive explanation of sectivity: the Intersectivity Axiom Schema holds when a certain frame composition is not immediately available in the linguistic and/or cultural competence of the speaker, or when the frame evoked by the type is very general. The second one makes sense, because when we refer to something with a very general type, we are probably referring to it in a global way, so that we cannot tell much about frame composition. On the contrary, when frame composition is typical, and the frame evoked by the type is not too general, the Intersectivity Axiom Schema does not hold. Here is how it works with respect to example (1). In *skilful surgeon* the frame evoked by *surgeon*, e.g. *Medical professionals* in FrameNet[2], can be conceptually composed with the frame evoked by *skilful*, i.e. *Expertise*[3], since expertise is a central feature of professionals. On the contrary, in *extroverted surgeon*, the *Medical professionals* frame does not typically compose with the frame evoked by *extroverted*, i.e. *Sociability*[4], since sociability is usually a secondary feature of scientific professionals. In examples (2) and (3) it is the generality of the frame that is involved. In (2), *dancer* is not a very general predicate for humans, while *fox* is the top level type for foxes. In

[2] https://framenet2.icsi.berkeley.edu/fnReports/data/frameIndex.xml?frame=Medical_professionals

[3] https://framenet2.icsi.berkeley.edu/fnReports/data/frameIndex.xml?frame=Expertise

[4] https://framenet2.icsi.berkeley.edu/fnReports/data/frameIndex.xml?frame=Sociability

(3), *scientist* is not a very general predicate for humans, while *woman* is one of the top level types. One may ask if non-rigidity [6] is actually behind these cases, since *dancer* and *scientist* are role-like predicates rather than types. This could be the case, since if we substitute *dancer* or *scientist* with e.g. *Japanese* or *New Yorker*, there is no sectivity switching, even though more specific types are used. Our initial hypothesis can then be reformulated as follows:

Cognitive Sectivity The Intersectivity Axiom Schema holds for any adjectival predicate (modifier) modifying a noun predicate (head) when: (a) the frames evoked by the modifier and the head are not composable in a conventional way, or (b) the frame evoked by the head is rigid. Framality will be studied empirically in future work, based on the extended adjective sets evoking FrameNet frames.

3. Related Work

A formal semantics for adjective sectivity interpretation has been often proposed in linguistic literature, including at least Bolinger's [7] and Vendler's [8] seminal work, the reference Montagovian approaches [9][4][10][2], and lexical semantics approaches along the lines of *generative lexicon* [11], and the Mikrokosmos ontology [12]. Partee [10] summarizes the mainstream Montagovian views on adjective semantics. The simplest rule was given by Montague [9]: *The denotation of an adjective phrase is always a function from properties to properties*, and any further constraint, e.g. sectivity, cannot be imposed to all adjectives, specially because of the existence of adjectives like *false, ostensible, alleged*. Kamp and Partee [13] provide a hierarchy and a set of constraints to make sense of special cases and contextual effects, except for intersective privative adjectives (cf. Sect.s 2 and 4. What we call here *non-extensional* are called in Montagovian literature *non-subsective*. Finally, Partee suggests that privative adjectives are better analyzed as subsective of a special kind, so giving room to the distinction that we propose here between intersective and subsective adjectives applied as privatives.

The formal linguistic literature provides very detailed accounts of adjective semantics; however, most work assumes adjectival *predicates*, trying to extend or correct the issues arising from the combination of adjective-evoked and noun-evoked predicates. The sectivity aspects of adjective semantics presented in Sect. 2 originate indeed from that assumption. While sectivity aspects need to be modelled, however, when we move from formal linguistics to formal ontology or even to semantic technologies in general, simpler approximate accounts are possible, and are advisable to start experimenting with real data extracted from large knowledge bases, textual corpora, etc.

Recently, a semantic-web-related contribution [14] has partly filled the gap between formal linguistics and formal ontology of adjectives, proposing a representation in the context of the so-called "ontology-lexicon interface", where lexical senses are represented as first-order semiotic entities, and are mapped to entities from existing ontologies. The distinction between intersective, subsective, and privative adjectives is annotated at the sense level, and their semantics in real ontologies is obtained via linking to ontology entities that are centred on either nouns or verbs, e.g. *Belgian* is represented in a Turtle serialization of OWL [15] as:

```
belgian:sense lemon:reference
  [ a owl:Restriction ; owl:onProperty dbpedia:nationality ;
      owl:hasValue dbpedia:Belgium ] .
```

which is (modulo namespaces) equivalent to the following description logic formula:

$$\texttt{reference(belgian,(}\exists\texttt{nationality.Belgium))} \tag{6}$$

This approach is reasonable, but estabishing the correct link is not trivial, even if the particular sense of the word is correctly disambiguated. For example, there is only one word sense for *Belgian* in WordNet, and it encompasses either nationality senses, or some broad sense of "origin": e.g. in *Belgian sole meunière Belgian* refers to the claimed origin of the recipe. It is therefore inappropriate to find a specific lexical mapping to existing ontologies for the open-ended set of senses that can be retrieved in discourse. It is certainly reasonable to do it for specific contexts, as the authors show in their evaluation with reference to question answering. Another limit is that the authors do not provide a specific solution for subsective adjectives, while for privative adjectives suggest a second-order representation that is outside of OWL expressivity. On the contrary, the representation proposed in this paper addresses all types of sectivity in a lightweight and reusable way. The main contribution of [14] is about gradable adjectives, and indeed a broader literature exists for the semantics of vague adjectives, where the seminal formal ontology work is Bennett's *relevant observable theory* [16].

In formal ontology, the closest literature is about *qualities* or *properties*, which are often expressed by adjectives in natural language. Qualities depart from the traditional treatment of adjective semantics in terms of predicates. Despite inter-translatability between qualities and adjectival predicates, qualities have the advantage to offer a distinct ontological type, and an integrated modelling of properties, property values, nuances and composition of properties, value metrics, etc., which distinguish most adjectival predicates from e.g. noun-evoked predicates.

The DOLCE family of foundational ontologies, including at least DOLCE-Full [17], DOLCE-Core [18], and DOLCE-Zero [19] (an extended OWL version for the Semantic Web) formalize qualities as individuals, related to *regions* that may belong to *quality spaces*, typically characterized by a metric. *Quality types* are also possible in order to assemble similar individual qualities.

For example, in *The Amazon River is wide* there is a predicative usage of *wide* without any explicit type assignment. In e.g. DOLCE-Zero, this gives room to the axiom:

$$\texttt{hasQuality}(AmazonRiver, wide_{AmazonRiver}) \tag{7}$$

where the individual quality $wide_{AmazonRiver}$ is inherent in *AmazonRiver*: it is the very way the Amazon River is wide (possibly vague, with additional idiosyncratic properties, etc.[5]). In addition, we may associate a *region* to this individual quality (e.g. a value in meters, possibly averaged from different widths in the course of the Amazon), and a *quality space* of lengths for that region. However, it is often hard to make a strong commitment to individual qualities and their values within regions of a metric space, therefore a generic property can be represented as a *quality type*, which can be further interpreted as an individual quality, or associated with a region, when more knowledge is available. The resulting axiom would be indeed underdetermined:

$$\texttt{hasQuality}(AmazonRiver, wide) \tag{8}$$

[5] Cf. also the literature on *moments* [20], *trope theory* [21], and local qualities [22]

4. Ontological and Representational Alternatives

The literature offers several different modelling approaches for adjective semantics, ranging from first-order to higher-order logics, from lambda calculus to modal logic, including a recent proposal for a description logic solution in the context of the ontology-lexicon interface for the Semantic Web. We propose some representation patterns and an ontology (see Sect. 5) enabling a simple formal modelling of adjectival predicates in discourse representation. The patterns are tailored to different sectivity types, and typically have multiple solutions. Some solutions are lightweight, ready to be used for knowledge graph[6] generation, others are more expressive, aimed at tasks requiring more precision. The solutions for each pattern are in principle inter-translatable.

4.1. Requirements

What we propose here is on one hand rigorous, on the other hand approximate and pragmatic, since the representation patterns and the ontology are used within an open knowledge extraction tool aimed at knowledge graph creation for linked data and the Semantic Web. The following is the list of requirements that we have taken into account for the creation of the patterns, and for the lexical ontology presented in Sect. 5:

- R1 Compatibility to state-of-the-art modelling of adjective sectivity distinctions
  - R1.1 Representation of as many as possible semantic constructs derivable from adjectival predicates
  - R1.2 Translatability to rich frame-based representation
- R2 Compatibility to quality semantics in formal ontology: adjectival predicates should be modelled as individual qualities whenever possible, otherwise as quality types
- R3 Compatibility to best practices in Linked Data and the Semantic Web
  - R3.1 RDF graphs with a basic OWL semantics to ensure consistency checking and automated inferences
  - R3.2 Simple schema, linked to existing ontologies when applicable
  - R3.3 Minimization of indirection: nodes in the graph should be as close as possible while respecting the intended semantics
  - R3.4 Generation of blank-node-free RDF graphs
  - R3.5 Avoid the creation of unnecessary entities
  - R3.6 Ontology design patterns [27] as RDF/OWL good modeling practices
- R4 Integration with open knowledge extraction patterns
  - R4.1 Smooth integration within Open Knowledge Extraction (OKE [28]) graphs
  - R4.2 A strict (as far as possible) correspondence between textual segments and entities in the formal model
  - R4.3 Clear and simple modelling of the Intersectivity Axiom Schema
  - R4.4 Optimal avoidance of inaccurate modelling for subsective, non-extensional, and privative adjectives
  - R4.4 Minimal need to use automated word-sense disambiguation (tool scalability requirement)
- R5 Integration with existing lexical resources encoding adjective semantics
  - R5.1 Consistency with the ontology-lexicon interface semantics typically used to model existing lexical resources, and performing inferences based on it
  - R5.2 Reuse of existing relations between adjectival words, senses, synsets, and frames

[6]The term *knowledge graph* is being used in scientific literature since a few years in order to generalize over a class of graph-structured knowledge bases containing factual data in form of (binary) relationships between identified entities. Curiously enough, the term lacks a formal definition, although there are two clear prototypes: RDF graphs [23] from Linked Data [24] and the Semantic Web, and proprietary graphs from Google [25], Bing [26], etc., which are built upon the ideas and resources from the Semantic Web.

R5.3 Reuse of linking between existing lexical resources

In summary, the considered requirements come from: formal semantics of natural language, formal ontology, linked data modelling patterns, open knowledge extraction practices, lexical semantics and resources.

4.2. Representation patterns

Intersective Pattern. The Intersective Pattern models the compositional semantics of adjective-noun phrases with intersective adjectives. With the quality semantics requirement (R2), and the Heim and Kratzer's (cit.) *Predicate Modification Rule*, the Intersectivity Axiom Schema (Ax. 2, requirement R4.3), is morphed as follows:

$$\texttt{C} \equiv \texttt{T} \sqcap \exists \texttt{hasQuality.M} \tag{9}$$

$$(semantics:)\ C^I = T^I \cap \{x \mid for\ \langle x,y\rangle \in hasQuality^I, y = M^I\} \tag{10}$$

Basically, Axiom 10 adds a quality kind M to the noun-evoked type T, and so defines C. The restriction $\exists hasQuality.M$ enables the inheritance of M to any individual typed by C, thus obtaining the same inferential behaviour of Axiom 2, but with M as an individual. The following is an example:[7]

$$\texttt{CanadianSurgeon} \equiv \texttt{Surgeon} \sqcap \exists \texttt{hasQuality.Canadian} \tag{11}$$

$$\texttt{CanadianSurgeon(Dr.\_Lara\_J.\_Williams)} \tag{12}$$

$$(11,12) \implies \texttt{Surgeon(Dr.\_Lara\_J.\_Williams)} \tag{13}$$

$$(11,12) \implies \texttt{hasQuality(Dr.\_Lara\_J.\_Williams,Canadian)} \tag{14}$$

Subsective Pattern. The Subsective Pattern models the compositional (taxonomical) semantics of adjective-noun phrases with subsective adjectives. Here we need to avoid the inheritance of quality kinds to individuals, since the quality works like an intensional modifier, and its extensional impact is mediated by the composed type. Notice that in OWL the *punning* mechanism (or any other meta-level sugar in other languages) allows to use a predicative constant as an individual constant: the alternative interpretation is contextually clear, and no higher-order reasoning is expected, *contra* second-order proposals for the representation of certain aspects of adjective semantics. Three solutions are proposed (S is any class of situations or n-ary reifications):

(intensional solution)		(trope solution)		(situational solution)	
$\texttt{C} \sqsubseteq \texttt{T}$	(15)	$\texttt{C} \sqsubseteq \texttt{T}$	(18)	$\texttt{C} \sqsubseteq \texttt{T}$	(22)
hasIntensionalQuality(C,M)	(16)	subQualityOf(M_as_T,M)	(19)	C(c)	(23)
C(c)	(17)	C(c)	(20)	S(s)	(24)
		hasQuality(c,M_as_T)	(21)	involvesIndividual(s,c)	(25)
				involvesType(s,T)	(26)
				involvesQuality(s,M)	(27)

Although the three solutions are in principle inter-translatable, they have pros and cons with respect to the requirements. The intensional solution is ideal for the requirements R3 and R4, but is a bit dubious in R1 and R2, e.g. qualities are expected to be properties of individuals (particulars), not predicates (universals), except in some cases such as privative adjectival predicates. The trope solution is probably the best for R2, because it preserves the individual applicability of qualities, but it requires the creation of

[7] Entities in examples are all real, and taken from the Web.

specialized qualities using a *qua-entity*-like approach [29], and their introduction through the `subQualityOf` relation. This modelling style is against the Okkam's razor principle used in linked data (R3.5), which suggest to avoid the introduction of extra entities when there is no urgent necessity to do it. It also breaks the correspondence between textual segments and the formal model (R4.2) maybe beyond necessity. The situation solution binds together individual, types, and qualities, and reifies the binding as a situation. This is very explicit, but it is probably against many requirements: new entities and relations to be introduced, much indirection in the resulting graph, etc. The following is an example with the intensional solution:

$$\texttt{SkilfulSurgeon} \sqsubseteq \texttt{Surgeon} \quad (28)$$

$$\texttt{hasIntensionalQuality(SkilfulSurgeon,Skilful)} \quad (29)$$

$$\texttt{SkilfulSurgeon(Dr.\_Lara\_J.\_Williams)} \quad (30)$$

$$(28,29) \implies \texttt{Surgeon(Dr.\_Lara\_J.\_Williams)} \quad (31)$$

For adjectives that have an ambiguous (intersective or subsective) sectivity, such as *good*, it is advisable to stick to the Subsective Pattern, since it is more conservative than the Intersective one.

Non-Extensional Pattern. The Non-Extensional Pattern models the semantics of adjective-noun phrases with non-extensional, typically modal, adjectives. Here the problem is that the quality is really intensional, and tends to alter the way in which the type can be used as a predicate. It is therefore natural to imagine an intensional solution, without taxonomy generation, but with a generic *associative* relation between the composed term, and the head type, which in principle needs to be completed in presence of additional knowledge. No inheritance is expected, contrary to previous patterns. We provide here (i.e., axioms 32 $\sim$ 34) an intensional solution for space reasons, but a situational solution is also applicable. With this pattern, similar observations apply as for the Subsective one (i.e., axioms 35–36).

(intensional solution)		(example)	
`hasModality(C,M)`	(32)	`associatedWith(AllegedThief,Thief)`	(35)
`associatedWith(C,T)`	(33)	`AllegedThief(Nour_Aljirudi)`	(36)
`C(c)`	(34)		

Subsective Privative Pattern. The Subsective Privative Pattern models the semantics of adjective-noun phrases with privative adjectives. The problem here is similar to the non-extensional one, and privative adjectives are typically considered higher-order, but they go beyond altering the usage of the modified type: they negate the core properties of the type. Barbara Partee [10] suggests, followed by Morzycki [2] a taxonomy extension solution, but a much simpler associative intensional solution, almost identical to intensional solution is also available.

(intensional solution)		(extended taxonomy solution)		(example)	
`hasIntensionalQuality(C,M)`	(37)	$\texttt{C} \sqsubseteq \texttt{T\_(broad)}$	(40)	$\texttt{FakeSphinx} \sqsubseteq \texttt{Sphinx\_(broad)}$	(44)
`associatedWith(C,T)`	(38)	$\texttt{T} \sqsubseteq \texttt{T\_(broad)}$	(41)	$\texttt{Sphinx} \sqsubseteq \texttt{Sphinx\_(broad)}$	(45)
`C(c)`	(39)	$\texttt{C} \setminus \texttt{T}$	(42)	$\texttt{FakeSphinx} \setminus \texttt{Sphinx}$	(46)
		`C(c)`	(43)	`FakeSphinx(GizaSphinx)`	(47)

The extended taxonomy solution gives us more insights into the relation between a privative predicate, and the deprived type, but requires the introduction of a new adhoc entity in the vocabulary, which is against modelling styles from R3. The intensional solution draws the usual caveats with R2 and possibly R1. A good example with the extended taxonomy solution is represented by the axioms 44 $\sim$ 47

Intersective Privative Pattern. The Intersective Privative Pattern models the semantics of adjective-noun phrases with adjectives that are not privative (they are actually intersective), but when associated with that type, they act as privatives. The solution here are structurally the same as for the Subsective Privative Pattern, but in this case the quality is inherited by the individual. For space reasons we only show an application example for this pattern:

$$\texttt{hasQuality(StoneLion,Stone)} \quad (48)$$

$$\texttt{StoneLion} \sqsubseteq \texttt{Lion\_broad)} \quad (49)$$

$$\texttt{Lion} \sqsubseteq \texttt{Lion\_broad)} \quad (50)$$

$$\texttt{StoneLion} \setminus \texttt{Lion} \quad (51)$$

$$\texttt{StoneLion(GizaSphinx)} \quad (52)$$

$$\texttt{hasQuality(GizaSphinx,Stone)} \quad (53)$$

5. Delving into Lexical Resources: an Adjective Ontology

Even after a good set of representational patterns has been designed, we were still far from applying adjectival predicates to automated knowledge extraction, since there was no resource classifying adjectives according to their sectivity. For this reason, we decided to create a novel resource by means of manual classification. However, there are 21,538 adjectives and 39,753 adjective senses in WordNet, and classifying each one requires a long effort even with crowdsourcing techniques. Therefore, we have designed a shortcut, following the lexical linked data requirements from R5.

5.1. Annotating frames, synsets and word senses

Firstly, we attempted to align the top-level of an adjectival taxonomy. Unfortunately, WordNet [30] does not contain even an informal taxonomy of adjective synsets (while there is a large one for noun synsets). The second attempt has been to search for any good criterion that aggregates adjectives according to their semantics. This time we have found two interesting resources: FrameNet [31] *frames*, which are associated with *lexical units* that are supposed to evoke them, and WordNet similarity relations, which are extensively used for adjective senses. By reusing the RDF/OWL version of FrameNet version 1.5, we have singled out the 296 frames that are evoked by 1786 adjectival lexical units. 196 of those lexical units may evoke more than one frame, then should be considered polysemous within FrameNet. With this resource, we have started annotating each frame with a sectivity (as well as other features), by running the Type Changing Test on each adjectival lexical unit supposed to evoke that frame. The hypothesis was that adjectives evoking a same frame also have the same sectivity. The hypothesis seemed confirmed after a sample testing, hence we annotated the full set of lexical units for the 296 frames. However, the coverage of adjectives was still scarce (1,786 is less than 10% of the full WordNet lexicon). Therefore, we have used another resource from the FrameBase project [8], which provides linking between FrameNet lexical units and WordNet word senses, and we have created a RDF encoding of this resource.[9] The 1,786 word senses belong

[8] http://framebase.org/files/data/other/framenet-wordnet-map.txt

[9] All the datasets described in this paper are available from http://www.ontologydesignpatterns.org/ont/adjectives/index.html

to 1,367 adjectival synsets in WordNet, and with this first annotated datasets, we have started broadening our coverage.

5.2. Extending adjectival frame coverage

The first inference that we wanted to draw was that since most synsets contain multiple word senses, if a synset inherits a sectivity annotation from a word sense that links to a lexical unit from a certain frame, then also the other word senses in that synset will have the same sectivity value. In order to draw the inference, we have integrated the RDF/OWL version of WordNet 3.0 updated from the pioneering 2.0 W3C version [32], with the new RDF dataset with annotated frames, lexical units, and word senses. In this way, the rich knowledge of WordNet ipso facto became functional to our task. By using a set of rules, implemented as SPARQL CONSTRUCT queries with entailment [33], we have propagated the frame sectivity values to WordNet synsets, and from there, to the word senses that were not linked to FrameNet lexical units. In addition, since word senses are linked by a similarity relation, we have also drawn the inference that similar adjective word senses would have the same sectivity value. A related result is that the automatically annotated word senses become new candidates for evoking one of the 296 frames, so enabling future empirical investigation of framality (cf. Sect. 2). A sample query is shown here for the propagation via similarity at synset level:

```
CONSTRUCT { ?syn1 adjsem:synsetSectivityBySimilarity ?sec }
WHERE {?f adjsem:frameSectivity ?sec .
  ?f skos:narrowerMatch ?syn . ?syn schema:similarTo ?syn1 }
```

The last heuristic rule that we have adopted is that of proper denominal adjectives, which derive from proper nouns: in this case, the semantics carried by the adjective is quite specific, and hence intersective. Following this rule, all 791 proper adjectives (featuring capitalization in WordNet) have been classified. After the application of all heuristics, 8,227 adjective word senses (corresponding to 6,678 unique adjectives) have received a sectivity value. The coverage is now 31% of adjectives.

5.3. Error analysis

Once a reasonable coverage has been reached, we have started evaluating the assignments by looking for conflicting assignments (intersective, subsective, ambiguous, non-extensional, privative) to a same synset: this denotes the presence of word senses that have different sectivity values in a same synset. Three types of error patterns were found, including:

- Wrong alignments of FrameNet lexical units to WordNet word senses in the FrameBase dataset. E.g., the synset `quiet-adjective-1` had been mapped erroneously to the frame `Sound_level`, while its alignment should be Volubility, since it refers to a behavioral property, rather than a sound one
- Missing word senses for FrameNet lexical unit, e.g. the lexical unit `covetous.1` may evoke either the frame `Desiring`, or `Emotion_directed`, which have different sectivity values. The issue is that no word sense exists for the emotion associated with desire, while FrameNet distinguishes desire from the associated emotion

- Wrong assignments of sectivity values to frames: in some cases, the annotators have not used a consistent decision making, so introducing incoherent values

Once the errors have been corrected, we came up with the following distribution of assignments for the 6,678 adjectives. Intersective: 3,272; subsective: 2,857; non-extensional: 175; subsective privative: 54; ambiguous: 320. An example of a monosemous adjective annotated with sectivity, and linked to WordNet and FrameNet data is shown here:

```
[]    a       adjsem:MonosemousAdjective ;
      adjsem:monoAdjectiveLabel "flabby"@en-US ;
      adjsem:frameOfMonoAdj frame:Body_description_holistic ;
      adjsem:synsetOfMonoAdj wn30instances:synset-soft-adjectivesatellite-15 ;
      adjsem:wordSenseOfMonoAdj wn30instances:wordsense-flabby-adjectivesatellite-1 ;
      adjsem:sectivity "inter" .
```

The lexical ontology with annotated adjectives has been associated with an algorithm (cf. Algorithm 1) to decide how to treat each quality extracted by FRED [1] in compositional contexts originated by ADJ+NOUN constructions. FRED is a knowledge extraction tool that reuses multiple NLP components in order to produce consistent, self-connected, and linked knowledge graphs from text. The previous representation pattern adopted by FRED for quality modelling was flat, and assigned any adjectival predicate as an individual quality, so generating obvious errors in all cases except intersective ones. With the integration of the new representation patterns and the adjective ontology, the errors dramatically lowered. In case an adjective has an ambiguous sectivity, or is not included in the list of annotated adjectives, we default to the Subsective Pattern, in order to maximize precision, even at the cost of losing recall. The preferred solutions for each pattern is always the intensional one because it is the closest to linked data and knowledge extractuion requirements.

5.4. Implementation

Algorithm 1 presents the pseudo-code of our algorithm. It has been implemented as a Java OSGi component on top of the current release of the software architecture of FRED[10]. The component is named FRED-Quant and it aims at activating the most appropriate solution for representing the adjectives recognised in the input sentence with respect to the our lexical ontology. The input of FRED-Quant consists of an RDF graph (i.e., g in Algorithm 1) that is generated by FRED from a natural language sentence. More details about the implementation of FRED and how it generates RDF graphs from natural language can be found in [1]. The ontology is stored in a triple store based on Jena TDB[11] and accessed/queried by using SPARQL [33]. The first operation (cf. line 2 in Algorithm 1) performed by FRED-Quant is to retrieve all the bigrams $< modifier, head >$ (i.e., *retrieveAdjectiveBigrams* procedure) where the modifier is tagged as an adjective. This is performed by means of a SPARQL query over *g*. The query aims at detecting such bigrams by exploting the part-of-speech tags available for each token in g. Then, the adjective is stored in the variable $adjective$ (line 4) and its nature (i.e., intersective, modal, privative, etc.) is verified by querying the lexical ontology (lines 5–12). Finally, according to the case that matches, the related solutions is applied. This results in an update of *g* that modifies the semantics of *g* by using the solutions described in Section 4.

[10] http://wit.istc.cnr.it/stlab-tools/fred
[11] https://jena.apache.org/documentation/tdb

Algorithm 1 FRED-Quant algorithm

```
procedure QUANT(g)
    bigrams ← retrieveAdjectiveBigrams(g)
    for each bigram ∈ bigrams do
        [adjective, head] ← getWords(bigram)
        if adjective ∈ intersective||adjective ∈ proper then
            applyExtensionalSolution(bigram, g)
        else if adjective ∈ modal then
            applyModalSolution(bigram, g)
        else if adjective ∈ privative then
            applyPrivativeSolution(bigram, g)
        else
            applyDefaultSolution(bigram, g)
        end if
    end for
end procedure
```

6. Conclusion

In this paper we analyse the main problems of adjective semantics: we provide a richer interpretation of the sectivity (meta-)property, and define a new property named *framality*. For both properties we propose a number of formal representation patterns that can be reused for representing knowledge automatically extracted from natural language text. In addition, we hypothesise that adjective sectivity can be determined by identifying the frame that it evokes, meaning that different adjectives evoking the same frame have the same sectivity behaviour. A manual inspection and testing of a sample of adjectives showed that the hypothesis is sound, hence after manually annotating a set of seeds, we exploited FrameNet lexical units and WordNet similarity relations for automatically generating an ontological resource of adjectives, classified based on their sectivity behaviour. Finally, we have integrated such ontology and the defined representation patterns into the FRED knowledge extraction tool for demonstrating an implementation of this method. Currently, we are conducting a more in depth investigation of adjective framality, which includes the use of crowdsourcing for assessing the soundness of our hypothesis on a large scale.

References

[1] V. Presutti, F. Draicchio, and A. Gangemi. Knowledge Extraction based on Discourse Representation Theory and Linguistic Frames. In *Proc. of EKAW 2012*, pages 114–129. Springer, 2012.

[2] M. Morzycki. *Modification*. Cambridge University Press, 2015.

[3] I. Heim and A. Kratzer. *Semantics in Generative Grammar*. Blackwell textbooks in linguistics. Blackwell publishers, Cambridge (Mass.), Oxford, 1998.

[4] P. Bankston. Modeling Nonintersective Adjectives Using Operator Logics. *The Review of Modern Logic*, 9(1–2):9–28, 2003.

[5] C. J. Fillmore. Frame semantics. In L. S. of Korea, editor, *Linguistics in the Morning Calm*, pages 111–137. Hanshin Publishing Co., 1982.

[6] N. Guarino and C. A. Welty. An Overview of OntoClean. In *Handbook on Ontologies*, pages 201–220. Springer, 2009.

[7] D. Bolinger. Adjectives in English: Attribution and Predication. *Lingua*, 18:1–34, 1967.

[8] Z. Vendler. *Adjectives and Nominalizations*. Number 5 in Papers on formal linguistics. Walter De Gruyter Inc., 1968.

[9] R. Montague. Universal Grammar. *Theoria*, 36(3):373–398, 1970.

[10] B. Partee. Compositionality and Coercion in Semantics: The Dynamics of Adjective Meaning. *Cognitive foundations of interpretation*, pages 145–161, 2007.

[11] J. Pustejovsky and S. Bergler. *Lexical Semantics and Knowledge Representation*. Springer, 1992.

[12] V. Raskin and S. Nirenburg. Lexical Semantics of Adjectives. *New Mexico State University, Computing Research Laboratory Technical Report, MCCS-95-288*, 1995.

[13] H. Kamp and B. Partee. Prototype theory and compositionality. *Cognition*, 57(2):129–191, 1995.

[14] J. P. McCrae, C. Unger, F. Quattri, and P. Cimiano. Modelling the semantics of adjectives in the ontology-lexicon interface. In *COLING 2014*, page 198, 2014.

[15] W. OWL Working Group. *OWL 2 Web Ontology Language: Document Overview*. W3C Recommendation, 27 October 2009. Available at http://www.w3.org/TR/owl2-overview/.

[16] B. Bennett. A Theory of Vague Adjectives Grounded in Relevant Observables. In C. Welty, P. Doherty, and J. Mylopoulos, editors, *Proc. of KR 2006*, pages 36–45, Lake District, UK, 2006. AAAI press.

[17] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider. Sweetening ontologies with DOLCE. In *Proc. of EKAW 2002*, pages 166–181, Sigüenza, Spain, 2002. Springer.

[18] S. Borgo and C. Masolo. Foundational Choices in DOLCE. In *Handbook on Ontologies*, pages 361–381. Springer, Berlin, Heidelberg, 2009.

[19] A. Gangemi. DOLCE-Zero Ontology. http://www.ontologydesignpatterns.org/ont/dul/d0.owl. Accessed: 2016-02-10.

[20] B. Smith and K. Mulligan. Framework for Formal Ontology. *Topoi*, 2(1):73–85, 1983.

[21] I. Johansson. Determinables as universals. *The Monist*, 83(1):101–121, 2000.

[22] N. Guarino. Local Qualities, Quality Fields, and Quality Patterns: A Preliminary Investigation. In *Proceedings of the Second Interdisciplinary Workshop The Shape of Things, Rio de Janeiro, Brazil, April 3-4, 2013*, pages 75–81, 2013.

[23] D. Wood, M. Lanthaler, and R. Cyganiak. *RDF 1.1 Concepts and Abstract Syntax*. W3C Recommendation, February 2014. Accessed: 2016-02-10.

[24] T. Heath and C. Bizer. *Linked Data: Evolving the Web into a Global Data Space*, volume 344. Morgan & Claypool Publishers, Synthesis Lectures on the Semantic Web edition, 2011.

[25] A. Singhal. Introducing the Knowledge Graph: Things, not Strings. http://googleblog.blogspot.com/2012/05/introducing-knowledge-graph-things-not.html. Accessed: 2016-02-10.

[26] R. Qian. Understand Your World with Bing. https://blogs.bing.com/search/2013/03/21/understand-your-world-with-bing/. Accessed: 2016-02-10.

[27] A. Gangemi and V. Presutti. Ontology Design Patterns. In S. Staab and R. Studer, editors, *Handbook of Ontologies*, International Handbooks on Information Systems. Springer, 2nd edition, 2009.

[28] V. Presutti, A. G. Nuzzolese, S. Consoli, A. Gangemi, and D. Reforgiato Recupero. From Hyperlinks to Semantic Web Properties using Open Knowledge Extraction. *Semantic Web Journal*, 2016.

[29] C. Masolo, L. Vieu, E. Bottazzi, C. Catenacci, R. Ferrario, A. Gangemi, and N. Guarino. Social Roles and their Descriptions. In *KR*, pages 267–277, 2004.

[30] G. A. Miller. WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11):39–41, 1995. DOI: 10.1145/219717.219748.

[31] C. F. Baker, C. J. Fillmore, and J. B. Lowe. The Berkeley FrameNet Project. In *Proc. of COLING 1998*, pages 86–90, Montreal, Quebec, Canada, 1998. ACL.

[32] M. van Assem, A. Gangemi, and G. Schreiber. Conversion of WordNet to a standard RDF/OWL representation. In *Proc. of LREC'06*, pages 237–242, Genova, Italy, 2006. ELRA.

[33] S. Harris, A. Seaborne, and E. Prud'hommeaux. SPARQL 1.1 Query Language. W3C recommendation, W3C, March 2013. https://www.w3.org/TR/sparql11-query/.