Glossa a journal of general linguistics

Andreou, Marios. 2017. Stereotype negation in Frame Semantics. *Glossa: a journal of general linguistics* 2(1): 79. 1–30, DOI: https://doi.org/10.5334/gjgl.293

RESEARCH

Stereotype negation in Frame Semantics

Marios Andreou
Heinrich Heine University Düsseldorf, DE
Marios.Andreou@uni-duesseldorf.de

In this study, I enquire into word-formation semantics and the way context specifies the reading of derived words. In particular, I use the apparatus of Frame Semantics and offer a treatment of lexical stereotype negation that is expressed in English by the prefixes *non-* (e.g. *nonanswer*, *noncolor*) and *un-* (e.g. *unpolitician*, *un-diva*). In order to account for this phenomenon, I introduce a formal treatment of lexical rules in Frame Semantics, making use of the formalism of attribute-value matrices. I also use corpus-extracted data to enquire into the way context impacts on scope properties. I motivate an analysis under which the "absence" of a characteristic of the base lexeme is treated as a change in the value of an attribute of the base lexeme. The treatment of lexical stereotype negation can advance our understanding of modification in word-formation semantics and lead to a more balanced analysis and understanding of all major categories.

Keywords: frame semantics; lexical rules; lexical negation; affixation; contextual coercion

1 Introduction

Stereotype negation is a type of lexical negation[1] in the domain of nominals. The term was introduced by Bauer et al. (2013: 365) to capture cases of negation in which the derived word (e.g. *non-person*, *unpolitician*) is generally a member of the category denoted by the base noun but lacks some of the qualities of the base noun. The derivative *unpolitician*, for instance, is a member of the category *politician* but lacks particular characteristics of this category. For example, the *unpolitician* may trail far behind other politicians in leadership ratings. In effect, the referent of the derived word is not a stereotypical exemplar of its category.

Although lexical negation by means of affixation has received considerable attention in literature (Jespersen 1917; Zimmer 1964; Marchand 1969; Funk 1986; Horn 1989/2001, 2002; Bauer & Huddleston 2002; Lieber 2004; Plag 2004; Kjellmer 2005; Hamawand 2009; Bauer et al. 2013), not all types of lexical negation have been studied to the same extent. In particular, although "contrary" and "contradictory" types of negation are rather well studied, other types, such as stereotype negation have received less attention.

The treatment of lexical negation from a formal perspective is also a desideratum (some approaches within the Lexical Semantic Framework for morphology can be found in (Lieber 2004; Andreou 2015). The fact that from a formal perspective, lexical negation, in general, and stereotype negation, in particular, are understudied, is partly due to the fact that not much attention has been paid to (a) the study of the semantics of nominals and

[1] The term that is usually used for the function of negative prefixes such as *un-* is "affixal negation" (Zimmer 1964). In this study, I will use the term "lexical negation" instead of "affixal negation" to refer to the contribution of *un-* and *non-*, because affixes can be used to function as sentential negators in several languages, as for example in Turkish. Given that "sentential (or clausal) negation" can be expressed by means of affixation in many languages, the distinction between lexical and sentential negation should not be based on whether the negator is an affix or a word (Dahl 2010).

(b) the study of the interaction between the semantics of derivation and the semantics of the base lexeme.

In particular, lexical-semantic studies have largely focused on the study of verbs and little attention has been directed towards the study of the representation of nominals (Jackendoff 1990; Levin 1993). Verbs have an argument structure which lends itself readily to a lexical semantic analysis, whereas nominals come with few if any relational arguments (e.g. *father of X*), and this renders the study of the internal structure of nominals very difficult (Löbner 2013: 313). The work of scholars such as Löbner (1985, 2011), Pustejovsky (1995), and Lieber (2004) has, nevertheless, shown that nominals, like events, can be studied in their own right.

Another area of semantics that calls for research, is the interface between word formation and lexical semantics. That is, the study of the interaction between the semantics of word formation processes and the semantics of the base lexeme. Despite recent attempts to formalize this interaction, as for example in the seminal work of Lieber (2004, 2007, 2010, 2015, 2016) and Jackendoff (2009), the way modification works in derivation merits further investigation.

In the present study, I aim to fill a gap in the study of nominals and the morphology–lexical-semantics interface by offering a treatment of lexical stereotype negation (e.g. *nonanswer*, *un-pool*). A treatment of stereotype negation will allow us to enquire into issues pertaining to how best to account for modification in word formation, how best to capture generalizations in the lexicon, what it means to assume that a derived noun lacks certain characteristics of the base noun, which types of attributes are in the scope of stereotype negation, and how the context specifies the reading of derived words.

In order to formalize the analysis I use the apparatus of Frame Semantics (Petersen 2007; Kallmeyer & Osswald 2013; Löbner 2014)[2] and propose to analyze complex words in terms of lexical rules that capture generalizations in the lexicon. In addition, in the spirit of the analysis offered by Lieber (2016), I use corpus-extracted data to enquire into the way context specifies the reading of derived words.

In the following, I give a brief overview of the literature on lexical negation and present the prefixes *non-* and *un-* which will concern us in the present study (Section 2). In Section 3, I present the apparatus of Frame Semantics and in Section 4 I delve more deeply into the analysis of word formation in Frame Semantics. I argue in favor of the introduction of lexical rules into the inventory of Frame Semantics and analyze stereotype negation. In 4.2, I focus my attention on the types of attributes stereotype negation has scope over. Section 5 concludes the paper.

2 Lexical negation

The most recent classification of prefixal lexical negation in English is the one provided in Bauer et al. (2013). In particular, Bauer et al. (2013) show that prefixal lexical negation in English covers a range of eight readings as illustrated in Figure 1. This range of negative readings is expressed by a number of affixes which can be polysemous, that is, they can give rise to multiple readings (for the distribution of these readings per affix and base category see Bauer et al. 2013: 367).

[2] Frames have been used by several scholars to model linguistic phenomena (for an overview see Lehrer & Kittay 1992). Frames, for example, figure in works on Lexical Functional Grammar (Bresnan 2001) and Head-Driven Phrase Structure Grammar (Pollard & Sag 1994). More recently, Sag (2012) uses a version of Frame Semantics and Minimal Recursion Semantics (Copestake et al. 2005). Fillmore's frames (Fillmore 1982) are used in the FrameNet project (Fillmore & Baker 2010). In this paper, I will use Frames as defined in the work of Petersen (2007), Kallmeyer & Osswald (2013), Löbner (2013, 2014, 2015), and Petersen & Gamerschlag (2014).

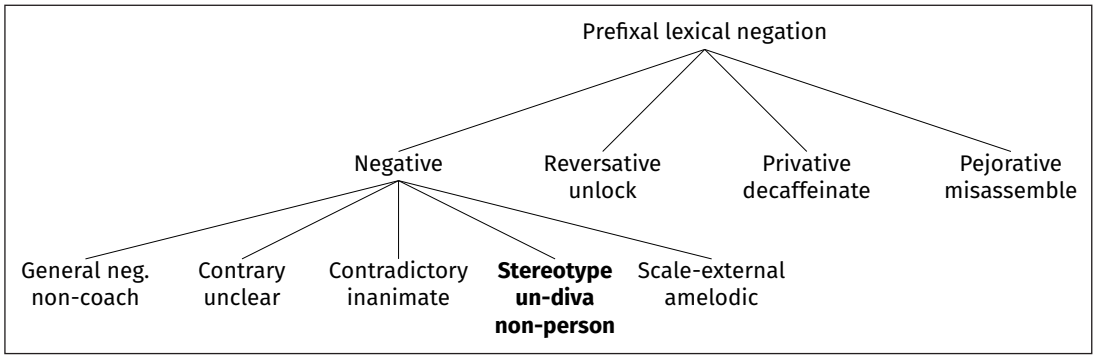

Figure 1: Distribution of readings in prefixal lexical negation (based on Bauer et al. 2013: Chapter 17).

In their classification, Bauer et al. (2013) propose four types of negativity, i.e. reversative ('reverse action of X-ing'), privative ('without X/remove X from'), pejorative ('do X wrongly'), and negative ('not X'). Reversative covers prefixed verbs such as *unlock* in which the derived verb describes a reversal of the action. Privation, which also covers removal, manifests itself in verbs such as *decaffeinate* in which the meaning 'depriving of or removing the thing described by the nominal base' is salient. Pejorative prefixes, such as *mis-* contribute a negative evaluation as in *misassemble* 'to assemble incorrectly'.

The negative type is split into five subtypes, namely general negative type, contrary, contradictory, scale-external, and stereotype negation.

The general negative reading is manifested in cases such as *non-coach* in which the derived lexeme denotes someone who is not a coach.

The distinction between contrary and contradictory negation is based on the characteristics of the adjective that serves as the base for the derivation. In particular, in contrary negation, "P" and "not-P" can be false at the same time since there can be a middle state between the two as exemplified by the pair *clear-unclear*. Although *clear* and *unclear* cannot be true at the same time, they can be simultaneously false since between the two, which are considered as terminal points on a gradable scale, there can be intermediate states. That is, something can be neither *clear* nor *unclear*.

Contradictory meanings, in contrast, exclude any intermediate states. To adduce an example, one can be either *animate* or *inanimate*; there is no middle state between the two (for more on this issue see Horn 1989/2001).

Scale-external negation is evident in cases in which what is denoted is the "complete irrelevance of the scale or polar opposition in question" (Bauer et al. 2013: 365). The adjective *amelodic*, for example, does not denote something melodic or unmelodic, but something for which the total absence of melody is relevant.

Let us now turn to stereotype negation, which concerns us in the present study. Bauer et al. (2013: 365) introduce this type of negation as follows:

> "In cases of stereotype negation, a noun is taken to denote a bundle of characteristics or qualities {x, y, z, ...}. When certain affixes are attached, what is negated is not the meaning of the noun as a whole, but a number of its semantic characteristics or qualities. The resulting derivative still generally denotes the same entity or something close to it, but one that is missing several key characteristics. In effect, the noun denotes a non-stereotypical exemplar of its category." (adapted from Bauer et al. 2013: 365)

Thus, in this type of negation, the derived word is generally a member of the category denoted by the base noun and what is negated is a part of the qualities of the base noun. A *non-person*, for instance, is a member of the category *person* that lacks particular characteristics of this category. In effect, the referent of the derived word is not a stereotypical exemplar of its category.

Could stereotype negation be subsumed under privation? In Bauer et al. (2013), privation and stereotype negation are distinct types, in that privation is a type in its own right, whereas stereotype negation is a sub-type (i.e. it belongs to the negative type along with the general negative, contrary, contradictory, and scale external sub-types).

Following the Aristotelian notion of privation, we could, nevertheless, treat stereotype negation as a sub-type of the privation type of negation (for an extensive discussion on negation in Aristotle see Horn 1989/2001). In particular, Aristotle's system of negation is based on opposition between pairs of expressions. Aristotle defines privation in terms of the absence or presence of a property that is expected to be inherent to the expression.

> "*Privation* and *possession* are spoken of in connexion with the same thing, for example sight and blindness in connexion with the eye. To generalize, each of them is spoken of in connexion with whatever the possession naturally occurs in. We say that anything capable of receiving a possession is deprived of it when it is entirely absent from that which naturally has it, and absent at the time when it is natural for it to have it. For it is not what has not teeth that we call toothless, or what has not sight blind, but what has not got them at the time when it is natural for it to have them." (Aristotle, Categories 12a26–33, trans. 1963)

The sub-type of stereotype negation introduced by Bauer et al. (2013) shares a crucial property with Aristotelian privation. That is, both involve the notion of absence/lack of characteristics. In particular, derived lexemes that belong to the sub-type of stereotype negation lack particular characteristics of their respective category. Crucially, they lack characteristics that are expected "by nature" to be inherent to these items. But what exactly does it mean to assume that a characteristic is absent from an item?

Before we address this question, we turn our attention to the prefixes *non-* and *un-* which serve as stereotype negators in English.

2.1 The prefix non-

The prefix *non-* attaches productively to both adjectives and nouns. It is polysemous and produces general negative readings, contrary readings, contradictory readings, and stereotype negation readings.

Given the nature of adjectives, *non-* on adjectives derives both contrary readings, as in (1a) and contradictory readings as in (1b); data come from the Corpus of Contemporary American English (COCA) and the British National Corpus (BNC), respectively:

(1) a. Contrary reading (COCA MAG 1993)
Let me get a **nonsubstantial** matter off the table at the start.

b. Contradictory reading (BNC 1992)
In 1963, the BBC was allocated a second TV channel (BBC 2) and by the end of the decade there were also a large number of local, commercial, and **non-commercial**, radio stations.

The prefixed word *nonsubstantial* exemplifies the contrary reading because there can be intermediate states between *substantial* and *nonsubstantial*. The contradictory reading of *non-commercial*, however, excludes any intermediate states (unless by coercion).

With respect to nouns, *non-* produces general negative readings as in (2) and sterotype negation as in (3).

(2) General negative reading (BNC 1993)
To protect the 'genuine' third party, the consent of the member States had also to be brought to the notice of the other party to the treaty, a **non-member** of the organisation.

(3) Stereotype negation (COCA FIC 1993)
The farmhouse had once possibly been yellow, but it had faded to about the same **non-color** as Anna's bedspread.

In (2), *non-member* denotes a party that is simply not a member of the respective organization (i.e. general negative reading 'not X'). The word *non-color* in (3), however, has a different negative flavor. In particular, *non-color* is a member of the category denoted by the respective base noun; a *non-color* is a kind of *color*. The prefixed form does not, however, denote a stereotypical exemplar of its respective category.[3]

The distinction between general negative readings and stereotype negation is not always easy to draw out of context and specific readings are contextually determined. This distinction is based on whether *nonX* is a member of category *X* (i.e. stereotype negation) or not (i.e. general negative reading). Consider the following:

(4) a. General negative reading (COCA ACAD 2012)
They suggested that other libraries should consider the advantages and disadvantages of the four methods, which included the difficulty of matching university academic departments to Library of Congress classifications, the inclusion of non-English materials, the inclusion of **nonbook materials**, staff time required to compile the data, the type of binding, and the currency of the data when attempting to determine average prices for use in allocating funds to various disciplines.

b. Stereotype negation (COCA ACAD 2010)
In my writing workshops I often meet the equivalent writing hobbyists. They are people who are writing what I term "coffee-break books," simple-minded **nonbooks** that they turn out in short order.

In (4a), the noun *nonbook materials*[4] denotes all materials that are not books; i.e. general negative reading. This includes manuscripts, audio-visual material etc. In (4b), however,

[3] Other examples of *non-* as a stereotype negator from the Oxford English Dictionary (OED) include lexemes such as *nonperson*, i.e. 'a person who is regarded as non-existent or unimportant, or who is not considered as a person for purposes of entitlement to rights, etc.; an ignored, humiliated, or forgotten person', *nonanswer*, i.e. 'an answer that does not deserve to be called an answer; an inadequate or evasive answer', and *nonword*, i.e. 'an unrecorded or hitherto unused word; a word which has (or is regarded as having) no accepted meaning' (for more examples see Algeo 1971).

[4] In *nonbook materials*, the noun *nonbook* is used as a modifier. A theoretical account of the adjectival use of nouns is still an open issue in the literature and I remain agnostic with regard to the theoretical treatment of such cases. Possible solutions involve conversion, zero-derivation, and coercion that is triggered by the immediate syntactic context.

nonbook denotes a book that is simpleminded, and, thus, not a stereotypical exemplar of the category *book*; i.e. stereotype negation.

2.2 The prefix un-

The prefix *un-*, is also polysemous and attaches productively to adjectives, nouns, and verbs (for a data-rich analysis see Horn 2002, 2005). With respect to adjectives we find both contrary and contradictory readings as illustrated in (5):

(5) a. Contrary reading (COCA SPOK 2003)
They're way out to sea, and they have a perfect right to be there and it's important for them to be there. It's our only eyes in the somewhat **unfriendly** world.

b. Contradictory reading (COCA FIC 2015)
During that indelible time of torment, I was all on my own. And, I must say now, with no lack of pride, my result, my **undeniable** victory, was no less than a masterpiece.

In (5a), *friendly* and *unfriendly* can be false at the same time since there can be a middle state between the two, whereas in (5b) something is either *deniable* or *undeniable*.

On verbs, *un-* productively selects for verbs that imply non-permanent results and gives rise to reversative readings. For example, *unzip* and *unfold* in (6) describe a reversal of the action of *zipping* and *folding*, respectively.

(6) Reversative reading (COCA MAG 2015)
To expand the Pop Top, close the handle, **unzip** the top zipper, **unfold** the top expansion compartment, and **unzip** the top back zipper and pull the handle through the slot.

As noted by Horn (2002, 2005), *un-* derives two readings on nouns. In particular, *un-turkey* in (7a) is not a member of the category *turkey* and *un-Politician* in (7b) is a member of the category *politician*, but not a stereotypical exemplar of this category.

(7) a. COCA MAG (1999)
[...] there's even a replacement for the traditional centerpiece of your holiday dinner. Turtle Island Foods and Now and Zen each offer wheat gluten-based **un-turkeys**, complete with stuffing, gravy and a crispy brown soy skin.

b. COCA MAG (2000)
Bradley knows this better than anyone. Mr. Authentic, with the well-worn shoes and soft voice and goodness platform, who hired the ad agency to help package him as the unpackaged candidate, said he wanted a different kind of campaign, noble and high protein. But Gore wanted to do it the old-fashioned way. Bradley might survive as the **un-Politician** but only if he is willing to fight for it and, by extension, fight for us.

The example in (7b) refers to the 2000 Democratic presidential primaries in which Vice President Albert Gore Jr. faced U.S. Senator Bill Bradley. Bradley is considered to be the *un-Politician* and is contrasted to Gore who prefers doing things the old-fashioned way. With respect to some of Bradley's characteristics consider the following:

(8) http://edition.cnn.com
So why is Bradley, the "**unpolitician**," using two Senators in his bio ad? Maybe because polls show most voters still think of him first as a former basketball player–and because he trails far behind Bush and Vice President Al Gore in "leadership" ratings.

In this context, Bradley lacks the leadership skills which are usually associated with other politicians such as Bush and Al Gore. In effect, this characteristic renders him a non-stereotypical member of the category *politician*.

2.3 Evaluative nuances of meaning

In the relevant literature, the prefix *un-* is usually associated with a negative evaluation that is generally absent with lexemes derived by *non-* (Jespersen 1917; Zimmer 1964; Algeo 1971). Should the semantic representations of *un-* and *non-* capture this assumed difference? A closer analysis of our data does not support an analysis in which specific evaluative nuances of meaning must be part of the semantic make-up of either *un-* or *non-*.

Consider the following with respect to *un-*. The careful reader may have noticed that *unpolitician* carries different nuances of evaluation in (7b) and (8). It carries a positive evaluation in (7b), whereas in (8) it is used with a depreciatory nuance of meaning. In particular, the *unpolitician* Bradley in (7b) represents something new, wants a noble campaign, and fights for the people. The same person in (8), however, is the *unpoliticain* because he trails far behind Bush and Vice President Al Gore in leadership ratings (i.e. negative evaluation).

In a similar vein, whether *non*-derivatives carry a positive or a negative nuance of meaning, is context dependent. Consider for example the positive evaluation associated with *nonpolitician* in (9).

(9) COCA NEWS (2004)
I think Jon Corzine is one of the most well-liked and effective senators, "said Senator Charles E. Schumer, a Democratic colleague from New York." It goes No. 1 to his genuineness. He is sort of a **nonpolitician**, and that makes him a better politician.

This positive nuance of meaning contrasts with the clearly depreciatory nuance that is expressed by *nonbook* in (10).

(10) COCA ACAD (2010)
In my writing workshops I often meet the equivalent writing hobbyists. They are people who are writing what I term "coffee-break books," simpleminded **nonbooks** that they turn out in short order.

Thus, in our treatment of the prefixes *un-* and *non-*, we need not include information on evaluation.

In this section, we gave a brief overview of the negative readings expressed by *non-* and *un-*. In the following, we focus on the way we could formally analyze the phenomenon of stereotype negation. In order to model this phenomenon we will use the apparatus of Frame Semantics. As we will see in the following sections, Frame Semantics allows one to focus on a micro-level and be explicit with respect to the properties of the base lexeme negation has scope over.

3 Frame Semantics

In what follows, I give a brief overview of the way frames are used as formats for describing concepts, with an emphasis on word formation. Consider the following two hypotheses from Löbner (2014: 23–24):

> H1 The human cognitive system operates with a single general format of representations.
> H2 If the human cognitive system operates with one general format of representations, this format is essentially Barsalou frames.

These two hypotheses build on the work of Barsalou (1992a, b, 1999) and constitute the *Frame Hypothesis*. A frame is a recursive attribute-value structure that provides information about a referent and fulfills the following three uniqueness conditions (adapted from Löbner 2014: 27):

> *Unique frame referent*
> All attributes and subattributes recursively relate to one and the same referent. (For the graph representation, there is exactly one node, the central node, such that every other node can be reached from it via a chain of one or more attribute arcs).
> *Unique values*
> Attributes are partial functions: Every attribute assigns to every possible possessor exactly one value.
> *Unique attributes*
> Every attribute is applied to a given possessor in a frame structure only once. (All attributes assigned to a given possessor are mutually different).

Frames can be represented as either attribute-value matrices, as also used, for example, in Head-driven Phrase Structure Grammar (HPSG, Pollard & Sag 1994), or as directed graphs. In the latter formalization a frame is "a directed, connected graph with nodes labeled by types and arcs labeled by attributes" (Petersen & Osswald 2014: 248). Attributes are always functional, in that there cannot be two arcs labeled with the same attribute going out from one node. The central node is the reference node[5] and is marked by a double border; rectangular borders are used for arguments. Consider for example the partial frame of the concept >ball<[6] as a directed graph and as an attribute-value matrix:

Figure 2 informs us that the SHAPE of >ball< is round. The double border marks the central node that refers to the extension of the concept.

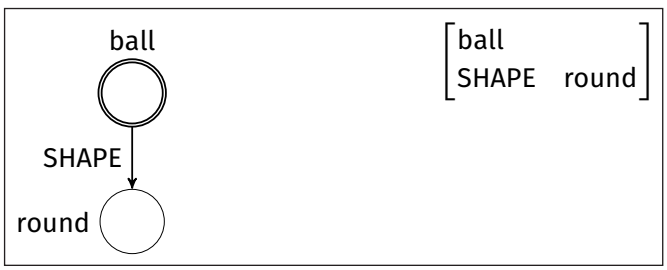

Figure 2: Partial frame for >ball< as a directed graph and as an attribute-value matrix.

[5] The reference node stands for the referential argument. In the case of nouns, for example, it stands for the so-called "R" argument that suggests "referential" and is involved in referential uses of NPs (Williams 1981; Wunderlich 2012).
[6] Concepts will be included in brackets > <.

3.1 Word formation in Frame Semantics

Word formation in Frame Semantics is generally treated in terms of referential shifts (see for example Löbner 2013 and Schulzek 2014 on *-er* and possessive compounds in German, and Kawaletz & Plag 2015 on English *-ment* nominalizations). The derived *walker* in Figure 3 serves as an illustrative example (from Löbner 2013: 312).

The concept >walk< has two attributes, namely AGENT and PATH as for example in *We walked to the station*. Thus, >walker< is formed by shifting the reference to the value of the attribute AGENT of >walk<. Observe that in accordance with bidirectional functionality, there is an attribute ACTIVITY that links the new referent back to the original referent node; a >walker< is engaged in a walking activity.

Kawaletz & Plag (2015) analyze *-ment* nominalizations in a similar manner. Consider, for example, the frame for the derived *bumfuzzlement* in Figure 4 (adapted from Kawaletz & Plag 2015: 312).

In their analysis, the verb *bumfuzzle* is a complex event of psychological causation and consists of two sub-events: a CAUSE and an EFFECT. The CAUSE is an activity and the

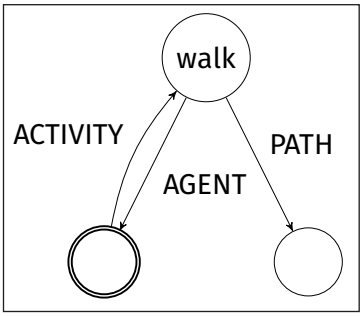

Figure 3: Frame for >walker<.

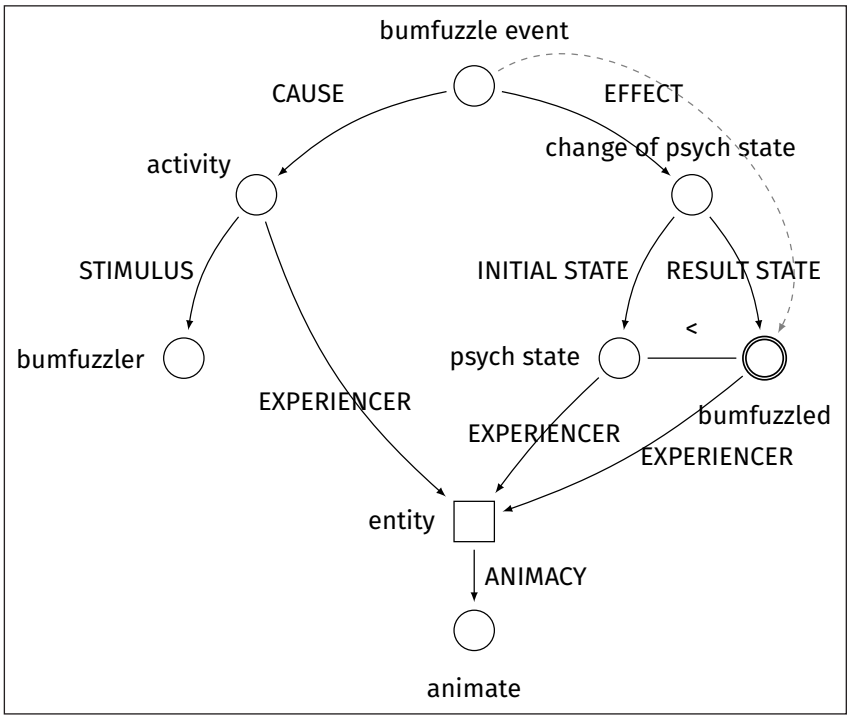

Figure 4: Partial frame for the nominalization *bumfuzzlement* in a RESULT STATE reading.

EFFECT is a change of psych state with an INITIAL STATE and a RESULT STATE. Thus, the RESULT STATE reading of *bumfuzzlement* is understood as a shift from the original referential node, i.e. *bumfuzzle event*, to the node *bumfuzzled*, which specifies the arc of RESULT STATE.

4 Stereotype negation in Frame Semantics

Lexical negation is a phenomenon in which information not already present in the frame of a concept seems to be added to it. This characteristic of lexical negation renders it different from both *-er* and *-ment* affixation. As we saw in Section 3, the concept >walker< is formed by shifting the reference from the original referential node to the value of the attribute AGENT of >walk<. In this case, the referent of >walker< is a participant in the >walk< event, and, thus, an argument already present in the frame of the base concept. In a similar vein, the result state reading of *bumfuzzlement* is a shift from the original referential node to the result state node. A treatment of negation in terms of metonymical shifts is not possible, however, for negation introduces semantic information that is not already part of the frame of the base concept.

The question now arises: How do we account for and model this kind of information? I propose to analyze word-formation semantics using rules that capture generalizations in the lexicon and model the interaction of affix and base semantics. In particular, I propose that stereotype negation can be analyzed in terms of a lexical rule that overrides the value of an attribute of the base lexeme. Which attribute(s) is/are affected is contextually determined.

Lexical rules have a long tradition in constraint-based models and have been used as a mechanism to reduce redundancy and to capture generalizations in the lexicon (Bresnan 1982; Pollard & Sag 1994; Briscoe & Copestake 1999; Bonami & Crysmann 2016). Such a rule for stereotype negation should not change the referential or categorial properties of the base word. That is, both the base and the derived word should share reference and belong to the same category.

I propose that lexical rules in Frame Semantics should be given in the form of an attribute-value matrix. Attribute-value matrices have been used by HPSG (Pollard & Sag 1994; Riehemann 1998; Koenig 1999) and other constraint-based models (see for example Bonami & Crysmann 2016 and literature therein) to capture morphological phenomena. In Frame Semantics, attribute-value matrices have also been used for syntactic and computational purposes (see for instance Kallmeyer & Osswald 2013; Osswald & Van Valin 2014). As I will show, the use of attribute-value structures proves to be very useful with respect to the analysis of lexical negation. In particular, it allows one to offer a detailed analysis of negation and to express scope. In addition, attribute-value matrices (contrary to graphs) make explicit reference to phonological and categorial features, and, thus, allow one to offer a treatment of negation in terms of lexical rules.

Another crucial notion in the apparatus of Frame Semantics which we need for our analysis is the notion of *type signature*. I assume that attributes and their values are given in a type signature which can be considered as an ontology which covers world knowledge (for more on typed feature structures see Carpenter 1992; Petersen & Gamerschlag 2014). According to Petersen & Gamerschlag (2014: 203–204) a type signature restricts the set of admissible frames, includes a hierarchy of the set of types, and states appropriateness conditions. These conditions declare the set of all admissible attributes for a lexeme of a certain type and the values these attributes take. Appropriateness conditions are inherited by subtypes. Consider, for example, the type signature in Figure 5.

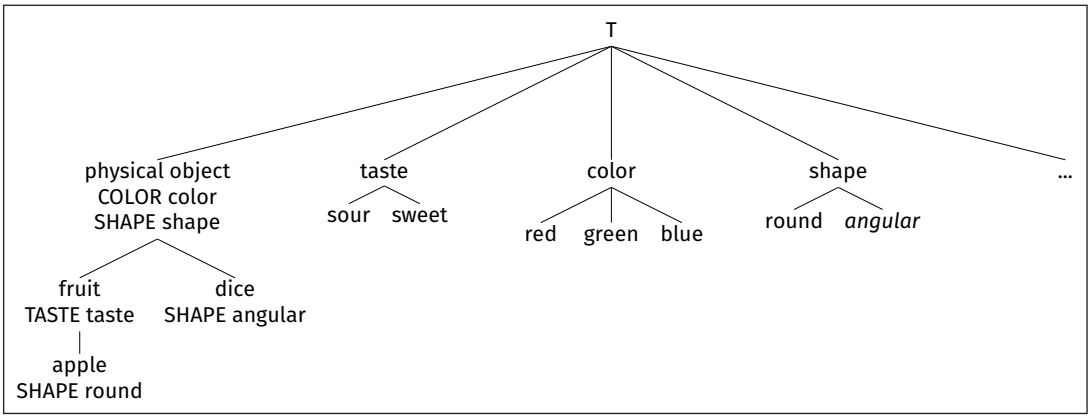

Figure 5: Example type signature (Petersen & Gamerschlag 2014: 204).

In this type signature, subtypes are given below supertypes. For example, *apple* is a *fruit*, which is itself a *physical object*. The node *physical object* meets two appropriateness conditions, that is, it is characterized by the attributes COLOR and SHAPE that have the values *color, red, green, blue* and *shape, round, angular*, respectively. According to the appropriateness conditions on *physical object*, TASTE does not attach to nodes of this type. Thus, not all *physical objects* have a taste. Given that appropriateness conditions are inherited and tighten by subtypes, *apple* inherits the appropriateness conditions on *fruit* and *physical object*. Thus, *apple* is characterized by the attributes TASTE, COLOR, and SHAPE. The value of SHAPE is *round* since subtypes not only inherit attributes from their supertypes, but also specify and tighten the value of inherited attributes. In a similar vein, *dice* inherits the attribute SHAPE from the node *physical object* and specifies the value of SHAPE as *angular*.

Let us now comment on the use of *color* as both an attribute label (i.e. COLOR) and a type label (i.e. *color*). In frames, this redundancy is attributed to the ontological status of attribute concepts. These functional concepts can be interpreted both *denotationally* and *relationally* (Guarino 1992). Thus, the denotational interpretation of color covers the set of all colors (i.e. type label *color*) and the relational interpretation covers the use of color as a functional attribute that assigns a particular color (e.g. *red*) to the referent of the frame (for more on the use of functional attributes see Löbner 2015).

4.1 Rule for stereotype negation

Figure 6 gives the rule for stereotype negation in the form of an attribute-value matrix.

The rule in Figure 6 gives a parallel representation of phonological information (PHON), morphosyntactic information (in particular, category, CAT), and semantic information (SEM) of both the derived lexeme and the morphological base (M-BASE). The M-BASE feature accounts for the internal structure of morphologically complex words and is equivalent to the morphological daughters notation (M-DTRS) used in Bonami & Crysmann (2016). Another important part of attribute-value matrices is structure sharing. Structure sharing is used to indicate that information in feature structures is identical. This is expressed by boxed numerals which are called *tags*, as for example $\boxed{1}$.

In Figure 6, the M-BASE has the phonology $\boxed{1}$, its category is noun, and its semantic information is specified by index (IND) and semantic frame (S-FRAME). IND identifies the referent of a lexeme. S-FRAME conveys two kinds of semantic information. First, it includes information on the referential properties (REF) of the M-BASE. Second, it includes functional attributes that assign values to the referent of the frame, in this case

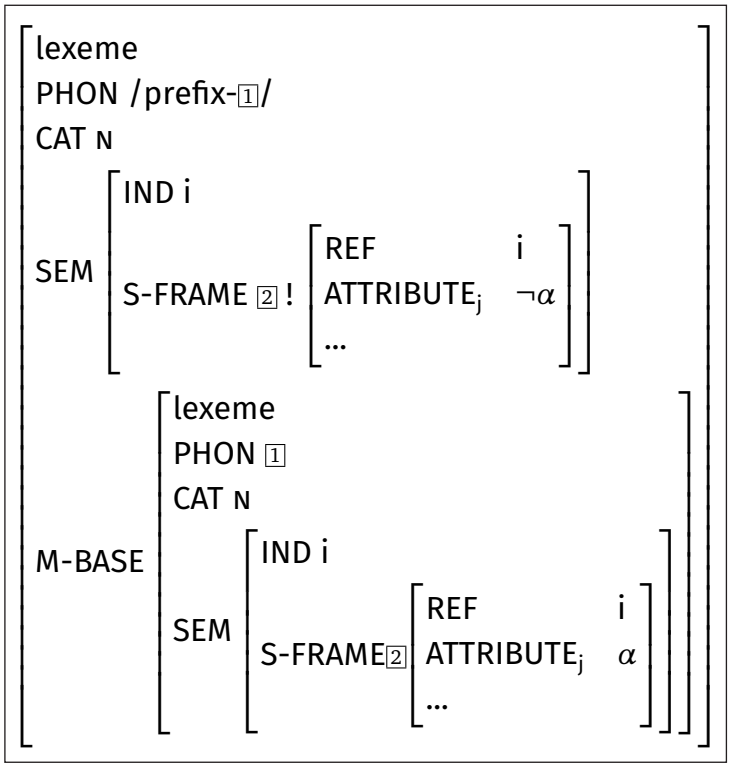

Figure 6: Rule for prefixal stereotype negation.

the referent of the M-BASE. In Figure 6, the M-BASE has an ATTRIBUTE$_j$ with the value α. The three dots in the S-FRAME indicate that there might be other attributes as well.

The derived lexeme is phonologically realized as /prefix-[1]/, where [1] is the phonology of the base lexeme. That is, the derived lexeme and the base lexeme share the phonological value [1]. The value of CAT is the same for both the derived and the base lexeme. That is, both are specified as nouns. The semantics (SEM) of the derived lexeme includes information on index (IND) and semantic frame (S-FRAME). The value *i* of IND and REF shows that the derived and base lexemes share reference. This is in accordance with the fact that *non-* and *un-* do not change the reference of the base lexeme.

The "S-FRAME [2] ![ATTRIBUTE$_j$ ¬α]" notation needs some explanation. First, the boxed numeral [2] shows that the value of S-FRAME of the derived lexeme must be identical to the S-FRAME part of the M-BASE. Second, the "!" notation, which is borrowed from Sag (2012: 125), informs us that the values of S-FRAME for the derived and base lexemes are identical except for the value of ATTRIBUTE$_j$. In particular, the value of ATTRIBUTE$_j$ is ¬α for the derived lexeme and α for the base lexeme.

The rule in Figure 6 is based on the type signature in Figure 7 which introduces a constraint on the values of attributes. In particular, the values α and ¬α of ATTRIBUTE$_j$ are not compatible with one another. From this it follows that negation of the value of an attribute alters/overrides the value of the attribute in question.

In a nutshell, the rule in Figure 6 provides a modeling of negation as a relation of opposition between pairs of items. In particular, in our rule we capture and model the way an item (i.e. the derivative) is opposed to another item (i.e. the base). In addition, we capture the generalization that *un-* and *non-* as stereotype negators attach to nouns and alter neither the category nor the reference of the base lexeme. In addition, this rule allows one to express scope. That is, the negation operator "¬" of the rule has scope over certain attributes of the base lexeme (i.e. ATTRIBUTE$_j$). It does not have scope over the whole base lexeme.

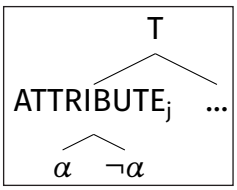

Figure 7: Type signature for ATTRIBUTE$_j$.

The override function of the rule for prefixal stereotype negation allows one to consider this phenomenon as an instance of coercion (Pustejovsky 1995, 2011). In particular, the negation operator takes scope over an implicit, yet conceptually salient attribute in the semantics of the morphological base and alters/overrides the value of that attribute. This view presupposes a dynamic relation between the various elements of an utterance and, most importantly, it presupposes that contextual elements can in fact interact with the elements of a lexeme and adjust its meaning (see also Lieber 2016). In particular, although the rule in Figure 6 allows us to model stereotype negation, it does not specify which particular attributes are affected by negation. As we will show in the next sections, it is the context which guides us to particular readings of derived lexemes. It is the context which determines which attributes negation has scope over and what values these attributes take.

The impact of context on lexical meaning has been the focus of recent work (among others Asher 2011; Lieber 2016). Lieber (2016) focuses specifically on the interplay between context and derived formations. She proposes to account for underspecification in nominalizations with an event reading (e.g. *The professor's examination of the student was thorough*) and a result reading (e.g. *The examination was two pages long*) in terms of contextual coercion. This means that underspecification in nominalizations is resolved within the larger context in which the nominalization occurs.

In the spirit of the analysis by Lieber (2016), in the following, we apply the rule in Figure 6 to data from *un-* and *non-*, and enquire into the way context can guide us to the interpretation of derived words. In particular, we will enquire into the way contextual information impacts on scope properties. A formalization of the incorporation of these scope properties in Frame Semantics using a constrained formal language is, nevertheless, not yet available and a topic of future research.

Although both *un-* and *non-* give rise to stereotype negation, in what follows, we do not collapse the two affixes. In particular, we treat them as two different affixes that are in competition (for more on competition in derivation see Lieber 2004; Bauer et al. 2013). We provide two different lexical rules, one for *un-* and one for *non-*, because a lexical rule must be fully specified with respect to phonological, morphological, and semantic information. Thus, the different values in the phonological part of *un-* and *non-* derivatives give rise to two distinct rules. It should also be mentioned that there could be other subtle meaning differences between the two prefixes but these differences did not show up in our data. In fact, any assumed subtle difference that could be included in the lexical rules for *non-* and *un-* would militate in favor of the idea that the two affixes must be treated separately. As a final point, it should be mentioned that in order to add a subtle difference in meaning to the core semantics of *non-* and *un-* as stereotype negators, one has to prove that the assumed difference is not to be attributed to other factors (e.g. contextual factors).

4.1.1 Modeling un-

Figure 8 gives the rule for *un-* as a stereotype negator.

This rule informs us that *un-* as a stereotype negator attaches to lexemes of CAT noun and alters the value of ATTRIBUTE$_j$. In effect, the resulting lexeme is not a stereotypical

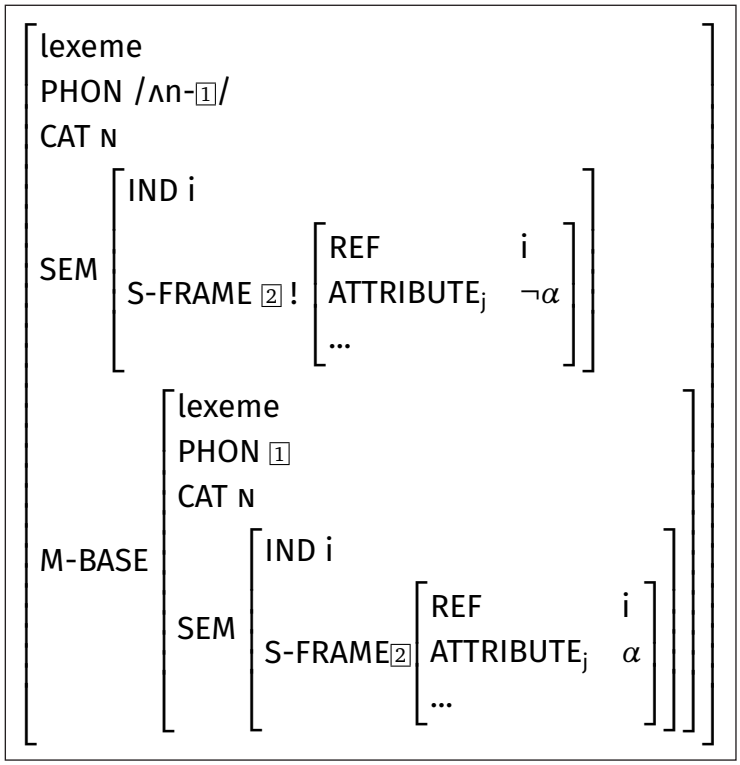

Figure 8: Rule for *un-* as a stereotype negator.

exemplar of the category denoted by the base. In addition, the resulting lexeme has the phonology /ʌn-1/ where 1 is the PHON of the morphological base.

Let us now delve more deeply into the way context feeds our rule. Consider first the following:

(11) COCA FIC (1993)
"What did you buy at the Galeries Lafayette?" [...] "Huh?" Half-asleep herself, Yael took a moment to grasp question. "Oh? You mean in Paris?" It all seemed to have hapned months ago. "Where then? In the Mitla Pass?" "Well, if you really want to know, some oo-la-la French **un-car**." "Ah! So? That's something to look forward to." He tried to die her as she wrapped the coat more closely around him, and struck his hand away."

In this example, there is no clear marker in the surrounding context that could help us identify why *un-car* is not a stereotypical exemplar of car. That is, there is no clear indication on the particular attributes negation has scope over and the values these attributes take. The only contextual element we could use is the modifier *French*, but it is not clear at all from the context what exactly *French* might refer to in this particular context. The participants in (11), may have some sort of shared knowledge or common ground which helps the addressee set particular values for particular attributes, but this is not something we can model by the context surrounding *un-car* in (11). In effect, the meaning of *un-car* remains largely underspecified. As modeled in Figure 9, the only values that the context sets are the values with respect to PHON.

In Figure 9, the values for PHON are set to /ʌnkɑː/ for the derived lexeme and /kɑː/ for the morphological base. The semantic part of the representation, however, is left underspecified. The semantic part allows us to model that *un-car* is not a stereotypical exemplar of car, but it does not spell out which attributes negation has scope over. In particular,

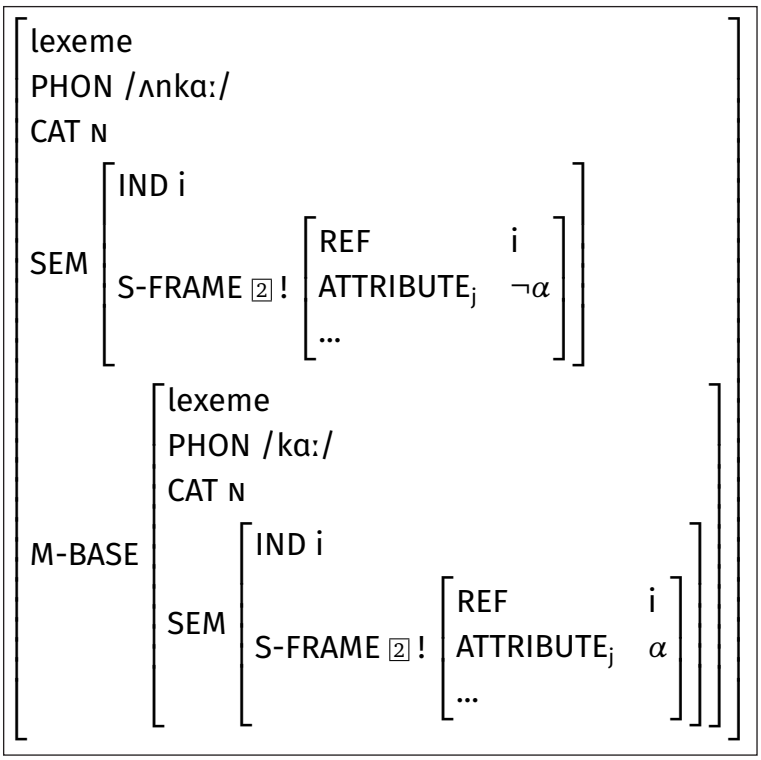

Figure 9: Attribute-value matrix for *un-car* (underspecified).

car and *un-car* differ with respect to the value of ATTRIBUTE$_j$; the value of ATTRIBUTE$_j$ is ¬α for the derived lexeme and α for the base lexeme. We cannot label ATTRIBUTE$_j$ and its respective values since in the surrounding context there is no clear marker we could base our analysis on. It should be mentioned that the absence of a clear marker which would flag the relevant attributes and their values is evident in several corpus-extracted examples.

Consider now the same prefixed word, i.e. *un-car*, in (12):

(12) COCA MAG (1997)
Here comes the "**Un-Car**," an old, multi-hued Honda pulled dogsled-style by three harnessed adults, two kids and a dog.

In contrast to (11), the context in (12) contains crucial information that helps us identify what renders the *Un-Car* a non-stereotypical exemplar of the category *car*. In particular, the important piece of information is given by "pulled dogsled-style by three harnessed adults, two kids and a dog".

How do we model this piece of information? In particular, how do we turn this into an attribute-value matrix? Such a task is fraught with complications, both methodological and expository. In particular, it is not always easy to come up with a label to name a functional attribute that assigns a value to a referent. In addition, labeling the possible range of values an attribute may take is not an easy task either.

The first step in our analysis is to identify the relevant attributes. In order to identify these attributes we need to base our analysis on inferences from the surrounding context in which *Un-Car* is found. In addition, we need to include a number of conceptually plausible functional attributes that are characteristic of the property described by the relevant piece of information. Based on the contextual information "pulled dogsled-style by three harnessed adults, two kids and a dog", we infer that the particular characteristic of *car* that is at issue in (12) is its POWER SOURCE.

The next step is to set the value of the functional attribute for the derived lexeme and the morphological base respectively. The POWER SOURCE of a stereotypical car is an *engine*. Thus, the value for POWER SOURCE in Figure 10 is set to *engine*. The *Un-Car* in (12), however, has a rather different POWER SOURCE. Its POWER SOURCE are actually some entities (i.e. "three harnessed adults, two kids and a dog") and not a standard engine.

An open issue is how we label the value for POWER SOURCE in the derived lexeme. There are two options available. The first option is to label the value as ¬engine (cf. ¬α). In effect, the value for the attribute POWER SOURCE will be an appropriate value but the value for the same attribute in the base. The second option is to label the value for POWER SOURCE based on inferences from the immediate context in which *Un-Car* is found. That is, the context in (12) fixes the value for POWER SOURCE to *entities*. Nothing important hinges on the specific choice of the labeling mechanism for the value of the relevant attribute in the derived lexeme. As we will see in this study, there are particular contexts in which we need to employ scalar attributes, such as attributes for color. In these cases, it will be very hard to choose a particular label for the value of the relevant scalar attribute. Thus, in order to avoid complication we shall use the ¬α strategy in our attribute-value matrices. Thus, in Figure 10 we fix the value for POWER SOURCE in the derived lexeme to ¬*engine*.

Note that the different values in *car* and *Un-Car* for the attribute BRAND are not related to stereotype negation. As we mentioned in Section 3, functional concepts can be interpreted both denotationally and relationally. Thus, the type label *brand* covers the set of all brands (denotational interpretation). The relational interpretation covers the use of BRAND as a functional attribute that assigns a particular brand to the referent of the frame. The particular value for BRAND in (12) is set to *Honda*.

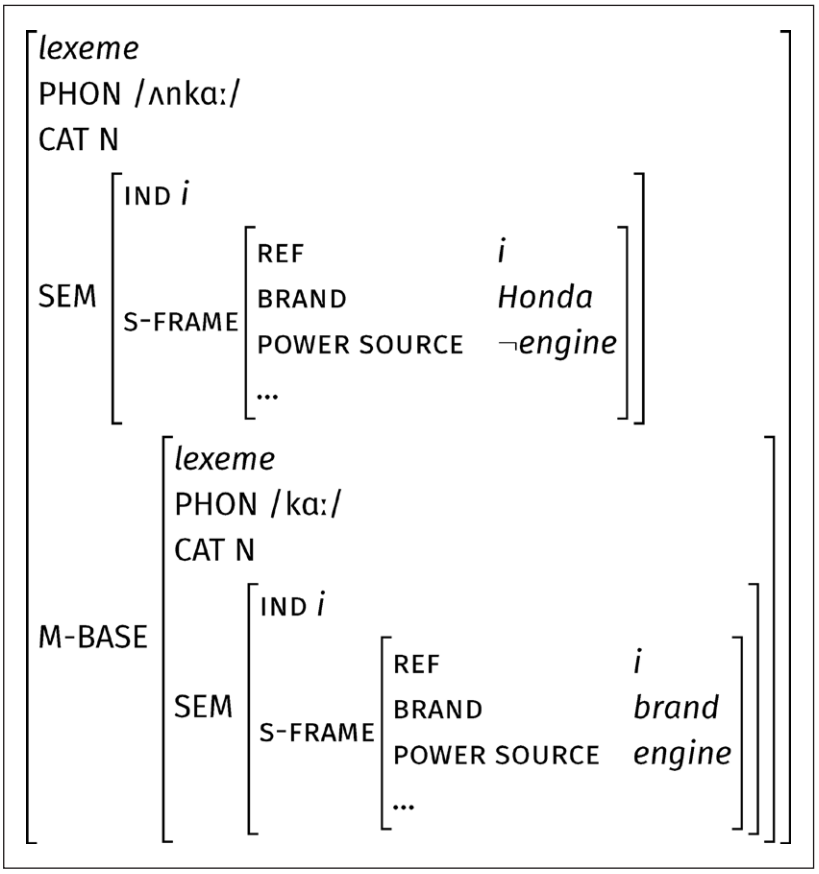

Figure 10: Attribute-value matrix for *Un-Car*.

In the previous example, the prefixed *Un-Car* denotes a car that is powered by entities. The *un-car* can be a non-stereotypical exemplar of *car* in other ways as well. To adduce an example, Mercedes Benz run an advertising campaign in the US for its Smart Car, the so-called "uncar". The campaign focused primarily on the size of the uncar (i.e. it is "unbig") and also on other characteristics, such as electric drive (i.e. the system which controls the motion of the electrical machines).

Consider also the following extract, in which emphasis is given on both the size and appearance of self-driving cars:

(13) http://www.spokesman.com
While research continues to create self-driving cars that can drive better than a human, there's also work to figure out the size and appearance of self-driving cars. [...] So what do you think? What will these vehicles of the near-future look like? Like the precious Google car? Golf cart-ish? Semi-truck-y? Like a worm? What should we call them? The **un-car**? The auto auto?

Let us now turn our attention to other prefixed formations. Consider the *un-diva* in (14):

(14) COCA SPOK (1994)
Dawn Upshaw has been called the **"un-diva"** of the opera world, often preferring to perform innovative, relatively obscure works that emphasize words over music in an informal style, often–imagine this–even chatting with an audience at recitals.

In this example, Dawn Upshaw is a diva who breaks down the stereotype for the category *diva*. In order to model the prefixed *un-diva*, we need to identify the relevant attributes negation has scope over and the values these attributes take. From the immediate context in (14) we infer that the relevant attributes are the STYLE of the *un-diva* (i.e. "preferring to perform [...] works in an informal style") and her APPROACHABILITY (i.e. "[...] even chatting with an audience at recitals"). In addition, there are three more attributes that relate to the works she prefers to perform. From the contextual information "innovative, relatively obscure works that emphasize words over music", we infer that the relevant attributes could be labeled INNOVATIVITY, OBSCURITY, and EMPHASIS, respectively. The attribute-value matrix in Figure 11 models the similarities and differences in values between the prefixed *un-diva* and its morphological base, i.e. *diva*.

The attribute-value matrix reads as follows: *un-diva* (like *diva*) is a female singer. In particular, she is an opera singer. In addition, *diva* has a *formal* STYLE of performing works, whereas the *un-diva* performs works in a rather informal STYLE. Finally, the *un-diva* is approachable, whereas the stereotypical *diva* is not. A note on the attribute-value matrix *singer* is in order. Given that frames are recursive structures, values for attributes are potentially complex frames themselves and, thus, can be further specified by attributes. In Figure 11, *singer* is a complex attribute-value frame that takes THEME as an attribute. The fact that frames are recursive attribute-value structures allows us to model the characteristics of the works the *un-diva* performs. That is, innovative, rather obscure works that do not emphasize music.

The analysis of the derived *un-diva* reveals an interesting characteristic of stereotype negation that has gone unnoticed in the literature. In particular, stereotype negation may have scope over attributes of sub-frames that are embedded into the frame for the original referent. In particular, the attributes INNOVATIVITY, OBSCURITY, and EMPHASIS do not assign properties to the referent of the *diva* frame. They provide information about the characteristics of the works the diva performs.

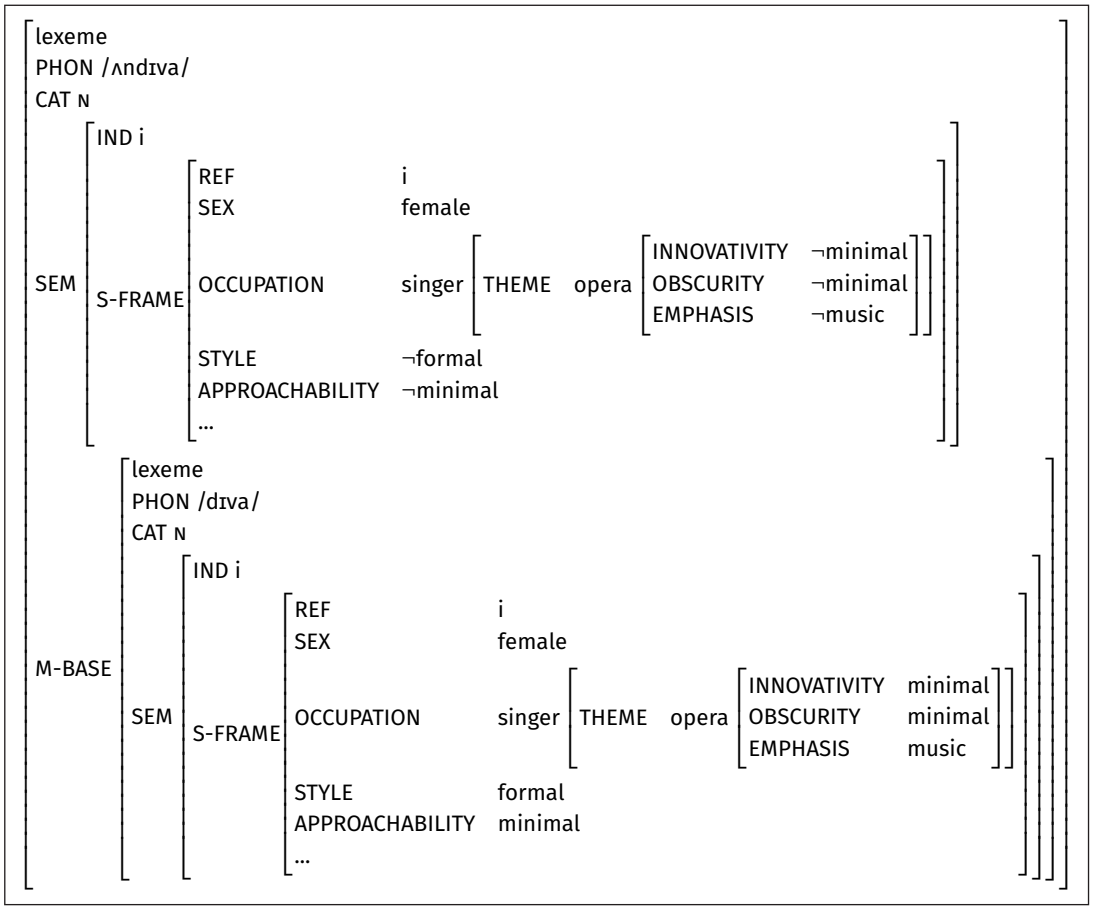

Figure 11: Attribute-value matrix for *un-diva*.

Consider now the prefixed *un-pools* in (15):

(15) COCA MAG (1999)
As for the **un-pools**, they tend to partner houses built by contemporary architects working in the stripped-down vernacular of modernism, where the water surface is just one more flat plane in an architectural assemblage of walls and roofs, masonry and glass. [...] In fact, pools are getting shallower around the country.

Based on the immediate context in which *un-pool* is found (i.e. "the water surface is just one more flat plane [...]. [...] pools are getting shallower[...]"), we infer that the *un-pool* differs significantly from a stereotypical *pool* with respect to its affordance.[7] Affordances of items describe the way we interact with them, i.e. the purpose or function of items. In Frame Semantics the respective label for the affordance attribute is FOR. The *un-pool* is a poorer member of the category denoted by the base lexeme, for it does not have the affordance of *recreation*. It serves an aesthetic function. In Figure 12, this is modeled as a difference in the value for the affordance attribute FOR of the morphological base and the derived lexeme.

In this section, we applied our rule for stereotype negation to words prefixed with *un-*. In the next section, we apply the same rule to *non-* derivatives.

[7] I am grateful to an anonymous reviewer for bringing this to my attention.

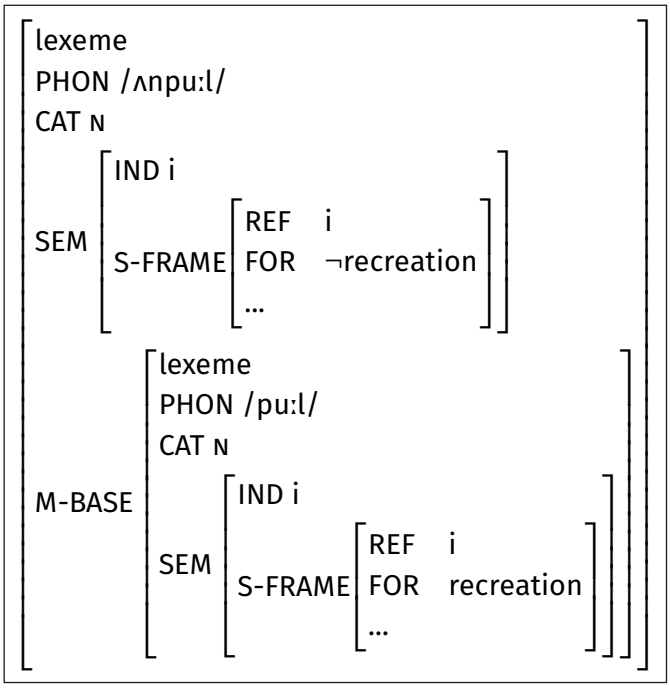

Figure 12: Attribute-value matrix for *un-pool*.

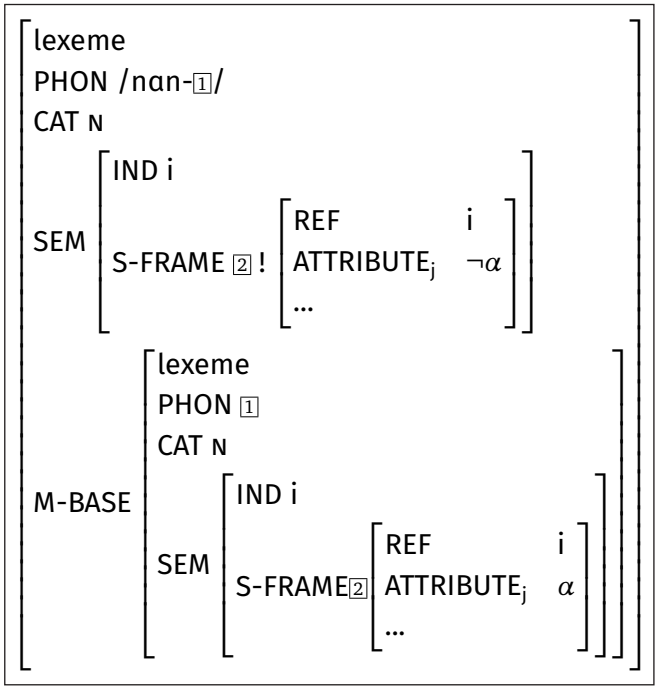

Figure 13: Rule for *non-* as a stereotype negator.

4.1.2 Modeling non-

Figure 13 gives the rule for *non-* as a stereotype negator.

This rule reads as follows: *non-* as a stereotype negator attaches to lexemes of CAT noun and alters the value of ATTRIBUTE$_j$ from α to ¬α. In effect, the resulting lexeme is not a stereotypical exemplar of the category denoted by the base. In addition, the resulting lexeme has the phonology /nɑn-1/ where 1 is the phonology of the morphological base.

Let us apply this rule to data and see how the immediate context in which *non-* words are found specifies their reading. Consider the use of *non-* as a stereotype negator in the following example from Bauer et al. (2013: 371):

(16) The man in the tweed suit wore his hair clipped short, in a crew cut. It was a flat metallic color, a **non-color**, like his eyes.

In this context, *non-color* is characterized as a color with a particular characteristic. It is a "flat metallic color". Based on this piece of contextual information we infer that *non-color* is a color which lacks the vibrancy usually associated with color. Thus, in our modeling of *non-color* we should include a number of conceptually plausible attributes that are characteristic of the property "vibrancy". I assume that this property is complex and that it is accounted for by the attributes INTENSITY and LUMINANCE. Figure 14 gives the attribute-value matrix for the derived lexeme *non-color* with the reading in (16).

Figure 14 reads as follows: the lexeme *non-color* is morphologically complex and its M-BASE is the lexeme *color*. The derived lexeme has the phonology /nɑnkʌlər/. *non-color* and *color* have the same values with respect to IND, REF, and CAT since *non-* changes neither the reference nor the category of the derived word. The derived *non-color* is a non-stereotypical exemplar of the category color since its values for INTENSITY and LUMINANCE differ from the values for the same attributes in the morphological base.

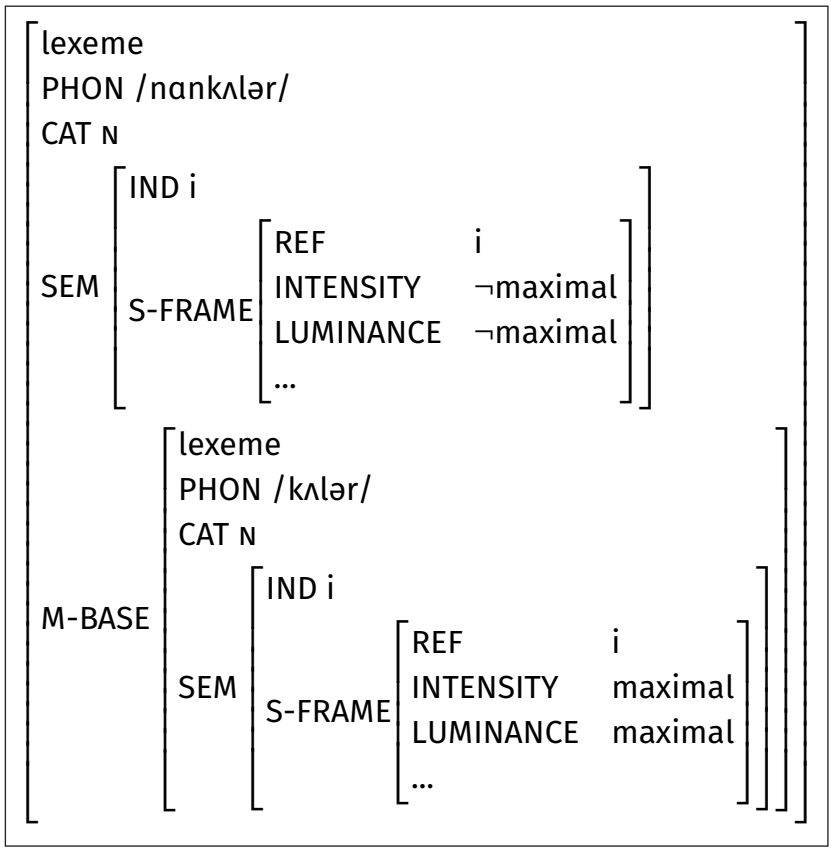

Figure 14: Attribute-value matrix for *non-color*.

Let us now consider two other cases of *non-color* that are exemplified in (17a) and (17b):

(17) a. COCA MAG (2002)
If the thoroughfares in this picturesque town really could talk, it's a sure bet that they would be buzzing about how the city's fashionable women are embracing the color white. And since that **noncolor** figures so prominently in this year's spring lines, women everywhere have plenty to choose from when it comes to finding just the "white" outfit.

b. COCA MAG (2010)
She wears predominantly black not only because it's the chosen **noncolor** of the working PR community but "also because I prefer to focus on the shape and the texture. Besides, when you travel as much as I do – I'm in a different city every week – everything must go with everything."

White in (17a) and black in (17b) are noncolor colors.[8] Let us examine whether what renders these colors non-stereotypical exemplars of the category *color* can be modeled in terms of the attributes INTENSITY and LUMINANCE. The immediate context in (17a) in which we find *noncolor* allows us to identify *noncolor* with the color white. The immediate context is "women are embracing the color white. And since that noncolor [...]". Based on this contextual information, we infer that the attribute that is at issue in (17a)

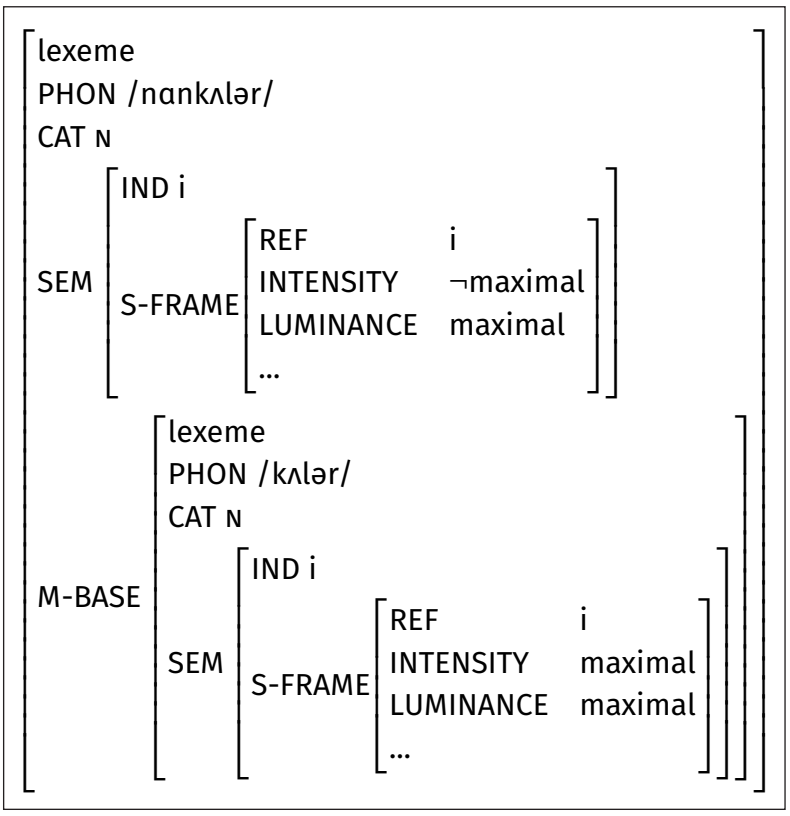

Figure 15: Attribute-value matrix for the *non-color* white.

[8] I borrow *noncolor color* from the following:

(i) COCA MAG (2006)
He is definitely offering a lot for the Jil Sander customer who has jumped ship in the past few years: classic Sander **noncolor colors**, the perfect white shirt, double-faced coats, great pants, and now even dresses!

is INTENSITY. Thus, as modeled in Figure 15, *non-color* with the reading 'white' is a color that has a different value with respect to INTENSITY.

In a similar manner, the contextual information "black not only because it's the chosen noncolor" in (17b) allows us to identify *noncolor* with the color black. From this information we infer that the attribute negation has scope over is LUMINANCE. Thus, *noncolor* with the reading 'black' is not a stereotypical exemplar of color since *color* and *noncolor* have different values for LUMINANCE. Figure 16 gives the attribute-value matrix for *non-color* with the reading 'black'.

Other examples of stereotype negation can be accounted for by the rule in Figure 6 in a similar manner. Consider *nonbook* in (18):

(18) COCA ACAD (2010)
In my writing workshops I often meet the equivalent writing hobbyists. They are people who are writing what I term "coffee-break books," simpleminded **nonbooks** that they turn out in short order.

Based on the immediate context "writing [...] coffee-break books, simpleminded nonbooks [...]" we identify *nonbook* as a particular kind of *book*. That is, as a "coffee-break" simpleminded book. From this contextual information we infer that the attribute of *book* that is at issue in (18) is the CONTENT of a book. This is modeled in the attribute-value matrix in Figure 17.

Figure 17 models the fact that *nonbooks* do not have the *complex* CONTENT which is stereotypically associated with books.

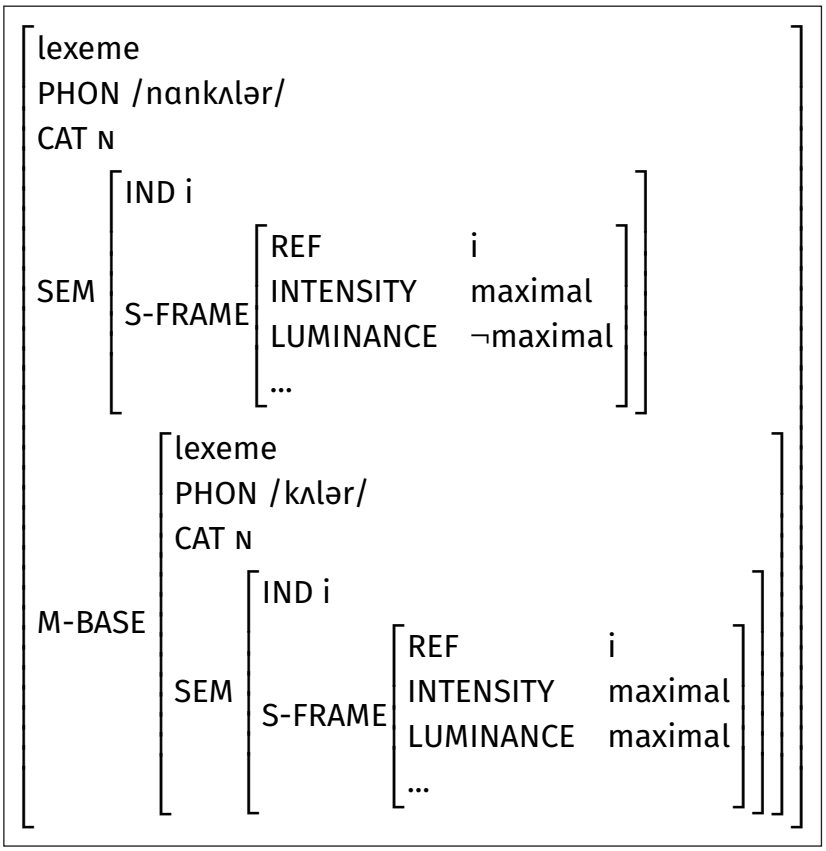

Figure 16: Attribute-value matrix for the *non-color* black.

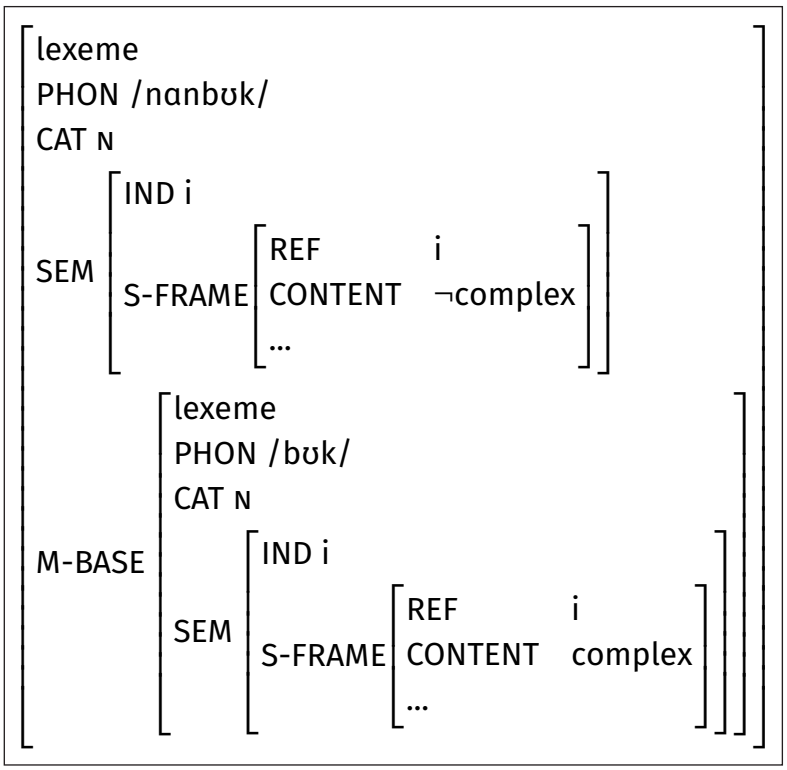

Figure 17: Attribute-value matrix for *nonbook*.

Consider now the prefixed *nonanswer* in (19):

(19) COCA NEWS (1995)
Like many artists who distrust discussion of their work, as though it will sap the energy needed to produce it, Van Sant is no interviewer's dream. He speaks affably about a number of subjects, but does so – in the space of this meeting, at least – at arm's length, in a facts-only fashion that disallows his personal feelings on any subject. # Asked if he was hurt by the "Cowgirls" reviews, for instance, Van Sant evades the question with a **nonanswer** about the nature of criticism.

In this example, the interviewed answers with a *nonanswer*. Based on the immediate context where *nonanswer* is found, we infer that a nonanswer does not provide relevant and adequate information. This can be modeled using the functional attribute RELEVANCE. In fact, the OED defines *nonanswer* as "an answer that does not deserve to be called an answer; an inadequate or evasive answer". Figure 18 models that *nonanswer* is a member of the category denoted by the base lexeme with which it shares form. The prefixed lexeme differs from its morphological base with respect to the value for the attribute RELEVANCE.

4.2 Stereotype negation and types of attributes

A final point to be addressed is the types of attributes stereotype negation has scope over. In particular, is there a constraint on the types of attributes stereotype negation has scope over?

Lexical semantic theorists have always tried to account for perceptual and encyclopedic properties of nouns. In the Lexical Semantic Framework of Lieber (2004), for example, aspects of meaning that are perceptual, cultural, and encyclopedic are encoded into the so-called *body* part of the lexical-semantic representation of nouns. In the Generative Lexicon of Pustejovsky (1995), these aspects of meaning are part of the *qualia stucture*

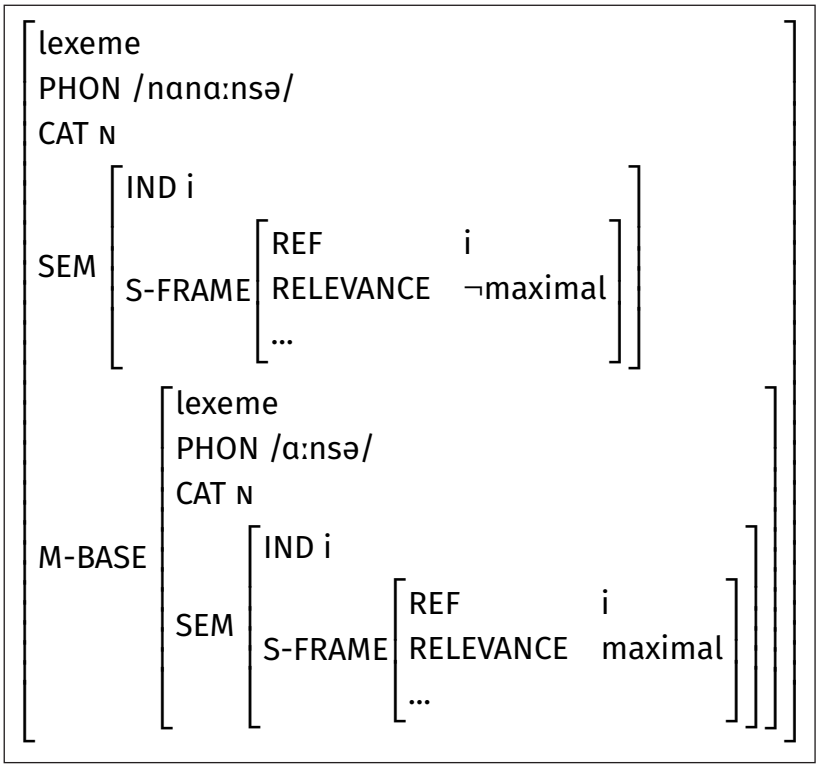

Figure 18: Attribute-value matrix for *nonanswer*.

level of representation. The *qualia structure* encodes information on what Pustejovsky (1995) calls the formal, constitutive, telic, and agentive aspects of the meaning of the lexical item.

Based on the types of values they assign, attributes in Frame Semantics are classified into four types: (a) part attributes, (b) property attributes, (c) event attributes, and (d) correlate attributes (Löbner 2013: 309). In what follows we provide an answer to the issue of which types of attributes are in the scope of stereotype negation.

First, part attributes provide a mereology of the referent. For example, the frame for *bicycle* will include a number of attributes such as WHEELS, PEDAL and so on and so forth. Some of these attributes will have their own mereology as well. The FRAME of a bicycle, for instance, includes the *top tube*, the *down tube*, the *seat tube*, the *chain stay* etc.

In our data, the *uncar* example in Figure 10 shows that stereotype negation can have scope over part attributes of the referent of the frame. In our example, it is the engine of the car that is affected by this type of negation.

Second, property attributes assign properties, usually of an abstract nature, to the referent of the frame. The COLOR or the WEIGHT of a bicycle are such property attritubes.

This is the most frequently attested type of attributes in our data, which is targeted by stereotype negation. Consider, for example, that INTENSITY and LUMINANCE in *non-color*, or the APPROACHABILITY of the *un-diva* are all property attributes stereotype negation has scope over.

Third, event attributes relate the referent to events and activities. In the *bicycle* frame, for example, there is an attribute that describes the affordance of *bicycle*. As we saw in the *un-pool* frame in Figure 12, this attribute has the label FOR. In the *bicycle* frame, bicycles are for *transportation*.

As evident in the *un-pool* example, stereotype negation may have scope over event attributes as well. What is in the scope of negation in the derived *un-pool*, is actually the *recreation* affordance of *pools*.

The last type of attributes is the so-called correlate-attributes type. Correlate attributes describe things of independent existence to which the referent of the frame is related. Such attributes are exemplified by the OWNER of the *bicycle* or the RESIDENCE of the *owner*. Consider the following example:

(20) COCA MAG (1993)
Perot would declare his candidacy the following Monday and, 10 days later, would follow up with a filmed prime-time speech disclosing his economic plan. His answer to the conventions would be a late-summer **unconvention**, a giant red, white and blue picnic, maybe in the Rose Bowl, maybe on the Mall in Washington; it would be a Norman Rockwell painting come to life.

In contrast to a stereotypical *convention*, Perot's *unconvention* features a number of characteristics, such as being "a giant red, white and blue picnic" and having no fixed place "maybe in the Rose Bowl, maybe on the Mall in Washington", which render it a non-stereotypical exemplar of the category *convention*. Let us focus on the PLACE attribute and examine how it differs from other attributes we have seen so far. This attribute is not a part attribute since it does not provide a mereology of the referent, it is not a property attribute since it does not assign a property to the referent, and it is not an event attribute since it does not relate the referent of the frame to an event. The PLACE attribute is different from other attributes we have examined so far, in that it describes a thing of independent existence (i.e. a place) to which the referent of the frame is related. Thus, it is a correlate attribute which has been targeted by stereotype negation. This shows that all types of attributes may be targeted by stereotype negation.

The proposed analysis highlights the importance of three aspects of frame theory. First, it shows that information on world knowledge has important ramifications for the analysis of lexical semantics. This is in accordance with previous work on decompositional models (see for example the work on the *qualia structure* by Pustejovsky 1995 and on the *encyclopedic body* by Lieber 2004). In particular, we saw that stereotype negation operates on the encyclopedic part of the base lexeme. Thus, a decompositional model that takes into account perceptual aspects of meaning allows one to focus on a micro-level and be much more explicit with respect to the properties of the base lexeme negation has scope over.

The second aspect is the distinction between functional attributes and the values they assign to referents in frame theory. This distinction allows us to answer a question we raised in Section 2 with respect to privation. In particular, what exactly does it mean to assume that a characteristic is absent from an item?

Our analysis suggests that stereotype negation does not have scope over the whole attribute. Rather, it has scope over the value of a given attribute. This has implications for the way stereotype negation works below the level of word, since the lack of a characteristic of the base lexeme is treated as a change in the value of an attribute of the base lexeme and not as absence of the attribute itself. This means that what is "absent" from an item is not the general property X, but the specific value Y for X, which is associated with the stereotypical exemplar of the category the item belongs to. Consider, for example, an analysis in which *non-* negates the attributes INTENSITY and LUMINANCE in *non-color*. Such an analysis would not predict the desired meaning, for *non-color* would be something to which INTENSITY and LUMINANCE are not relevant. The meaning of *non-color*, however, guides us to an analysis in which the attributes for vibrancy are still present in the derived word. They just have a different value. In a similar vein, the *un-pool* does not lack the general property of affordance. It has a different value for this property.

The third aspect is that frames are not mere lists of features but have internal structure. This structure consists of attributes and values that can be complex frames themselves as we saw in the case of *un-diva*. It is this particular structure that allows one to model the relation between the base lexeme and the prefixed lexeme and to analyze stereotype negation in terms of a rule that overrides the value of an attribute.

The proposed analysis also highlights the importance of context in determining the reading of derived words. In particular, the rule for stereotype negation in Figure 6 allows one to analyze this type of negation in terms of a rule that takes scope over attributes of the base lexeme. It is the context, however, that determines (a) which these attributes are and (b) what values these attributes take. For instance, the rule in Figure 6 allows us to model that the *uncar* is not a stereotypical exemplar of the category *car*, but it is the context that guides us to fix particular values for particular attributes.

5 Conclusion

The aim of the present paper was to enquire into the study of nominals and the interface between morphology and lexical semantics by offering a treatment of lexical stereotype negation. To this end, I analyzed corpus-extracted data using the framework of Frame Semantics.

The proposed analysis highlighted the importance of structured information on the micro-level of perceptual aspects of meaning. In particular, in Frame Semantics, this micro-level of analysis is expressed by functional attributes that assign properties to referents and the values these attributes take. This particular internal structure of frames allows one to tackle the issue of the "absence" of a characteristic from an item in a formal way. The answer we provided to this issue is that the "absence" of a characteristic of the base lexeme is treated as a change in the value of an attribute of the base lexeme and not as absence of the attribute (i.e. the general characteristic) itself.

The proposed analysis also highlighted the importance of context in specifying the reading of derived words. As we saw in our analysis of corpus-extracted data, the semantics of derived words is largely underspecified. Some derivatives, as for example *un-car* in (11), remain underspecified even when they are embedded in context. The reading of other derivatives, however, is specified by the context in which they occur.

As we have seen in our analysis, context serves three purposes. First, it accounts for evaluative nuances of meaning. This means that in our modeling of the semantic representation of *non-* and *un-*, there should not be a special evaluative part of meaning. Second, it disambiguates between general negative readings (i.e. "not X") and stereotype negation, as for example in the case of *nonbook* in (4). Third, the attributes negation has scope over and the values they take are fixed by using contextual information. In particular, although the rule in Figure 6 allows one to model stereotype negation, it does not specify which particular attributes are affected by negation. It is the context which guides us to the reading of derived lexemes since it determines which these attributes are and what values these attributes take. This view presupposes a dynamic relation between the context and the derived lexeme.

The introduction of a formal treatment of lexical rules in Frame Semantics offers a way to model word formation not in terms of conceptual metonymical shifts (such a treatment is not possible for negation), but in terms of fully specified rules and constraints that model the way the semantics of the affix interacts with the semantics of the base. As I hope to have shown in the present paper, the analysis of stereotype negation can advance our understanding of modification in word-formation semantics and lead to a more balanced analysis and understanding of all major categories.

Abbreviations

BNC = British National Corpus, CAT = category, COCA = Corpus of Contemporary American English, IND = index, M-BASE = morphological base, M-DTRS = morphological daughters, N = noun, OED = Oxford English Dictionary, PHON = phonology, REF = reference, SEM = semantics, S-FRAME = semantic frame, T = type signature

Acknowledgements

I would like to thank Rochelle Lieber, Sebastian Löbner, Laura Kallmeyer, Lea Kawaletz, Rainer Osswald, Ingo Plag, Marianna Gkiouleka, the editors of *Glossa*, and two anonymous reviewers for their comments and criticism. I would also like to thank Gero Kunter for his free corpus query tool *Coquery*. I gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft SFB 991 "The Structure of Representations in Language, Cognition, and Science" (Project: C08, "The semantics of derivational morphology: A frame-based approach").

Competing Interests

The author has no competing interests to declare.

References

Algeo, John. 1971. The voguish uses of non. *American Speech* 46(1/2). 87–105. DOI: https://doi.org/10.2307/3087991

Andreou, Marios. 2015. Lexical negation in lexical semantics: The prefixes in- and dis-. *Morphology* 25(4). 391–410. DOI: https://doi.org/10.1007/s11525-015-9266-z

Asher, Nicholas. 2011. *Lexical meaning in context: A web of words*. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511793936

Barsalou, Lawrence W. 1992a. *Cognitive psychology: An overview for cognitive sciences*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, Lawrence W. 1992b. Frames, concepts, and conceptual fields. In Adrienne Lehrer & Eva Feder Kittay (eds.), *Frames, fields and contrasts: New essays in semantic and lexical organization*, 21–74. Hillsdale, NJ: Lawrence Erlbaum Associates.

Barsalou, Lawrence W. 1999. Perceptual symbol systems. *The Behavioral and Brain Sciences* 22. 577–660. DOI: https://doi.org/10.1017/S0140525X99002149

Bauer, Laurie, Rochelle Lieber & Ingo Plag. 2013. *The Oxford reference guide to English morphology*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780198747062.001.0001

Bauer, Laurie & Rodney Huddleston. 2002. Lexical word-formation. In Rodney Huddleston & Geoffrey K. Pullum (eds.), *The Cambridge grammar of the English language*, 1621–1722. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/9781316423530.020

Bonami, Olivier & Berthold Crysmann. 2016. Morphology in constraint-based lexical approaches to grammar. In Andrew Hippisley & Gregory T. Stump (eds.), *Cambridge handbook of morphology*, 609–656. Cambridge: Cambridge University Press.

Bresnan, Joan. 1982. The passive in lexical theory. In Joan Bresnan (ed.), *The mental representation of grammatical relations*, 3–86. Cambridge, MA: MIT Press.

Bresnan, Joan. 2001. *Lexical-functional syntax*. Oxford: Blackwell.

Briscoe, Ted & Ann Copestake. 1999. Lexical rules in constraint-based grammars. *Computational Linguistics* 25(4). 487–526.

Carpenter, Bob. 1992. *The logic of typed feature structures*. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511530098

Copestake, Ann, Dan Flickinger, Carl J. Pollard & Ivan A. Sag. 2005. Minimal recursion semantics: An introduction. *Research on Language and Computation* 3(4). 281–332. DOI: https://doi.org/10.1007/s11168-006-6327-9

Dahl, Östen. 2010. Typology of negation. In Laurence R. Horn (ed.), *The expression of negation*, 9–38. Berlin: De Gruyter Mouton.

Fillmore, Charles J. 1982. Frame semantics. In The linguistic society of Korea (ed.), *Linguistics in the morning calm*, 111–137. Seoul: Hanshin Publishing Co.

Fillmore, Charles J. & Colin Baker. 2010. A frames approach to semantic analysis. In Bernd Heine & Heiko Narrog (eds.), *The Oxford handbook of linguistic analysis*, 313–340. Oxford: Oxford University Press.

Funk, Wolf-Peter. 1986. Towards a definition of semantic constraints on negative prefixation in English and German. In Dieter Kastovsky & Aleksander Szwedek (eds.), *Linguistics across historical and geographical boundaries: In honour of Jacek Fisiak on the occasion of his fiftieth birthday*, 877–889. Berlin: Mouton de Gruyter. DOI: https://doi.org/10.1515/9783110856132.877

Guarino, Nicola. 1992. Concepts, attributes and arbitrary relations: Some linguistic and ontological criteria for structuring knowledge bases. *Data and Knowledge Engineering* 8(3). 249–261. DOI: https://doi.org/10.1016/0169-023X(92)90025-7

Hamawand, Zeki. 2009. *The semantics of English negative prefixes*. London: Equinox.

Horn, Laurence R. 1989/2001. *A natural history of negation*. Chicago, IL: University of Chicago Press. Reprinted 2001 by CSLI Publications, Stanford, CA.

Horn, Laurence R. 2002. Uncovering the un-word: A study in lexical pragmatics. *Sophia Linguistica* 49. 1–64.

Horn, Lawrence R. 2005. An un-paper for the unsyntactician. In Salikoko S. Mufwene, Elaine J. Francis & Rebecca S. Wheeler (eds.), *Polymorphous linguistics: Jim McCawley's legacy*, 329–365. Cambridge, MA: MIT Press.

Jackendoff, Ray. 1990. *Semantic structures*. Cambridge, MA: MIT Press.

Jackendoff, Ray. 2009. Compounding in the parallel architecture and conceptual semantics. In Rochelle Lieber & Pavol Štekauer (eds.), *The Oxford handbook of compounding*, 105–128. Oxford: Oxford University Press.

Jespersen, Otto. 1917. *Negation in English and other languages*. Copenhagen: A. F. Host.

Kallmeyer, Laura & Rainer Osswald. 2013. Syntax-driven semantic frame composition in lexicalized tree adjoining grammars. *Journal of Language Modelling* 1(2). 267–330. DOI: https://doi.org/10.15398/jlm.v1i2.61

Kawaletz, Lea & Ingo Plag. 2015. Predicting the semantics of English nominalizations: A frame-based analysis of -ment suffixation. In Laurie Bauer, Lívia Körtvélyessy & Pavol Štekauer (eds.), *Semantics of complex words*, 289–319. Dordrecht: Springer.

Kjellmer, Göran. 2005. Negated adjectives in modern English: A corpusbased study. *Studia Neophilologica* 77(2). 156–170. DOI: https://doi.org/10.1080/00393270500355130

Koenig, Jean-Pierre. 1999. *Lexical relations*. Stanford, CA: CSLI Publications.

Lehrer, Adrienne J. & Eva Feder Kittay (eds.). 1992. *Frames, fields, and contrasts*. Hillsday, NJ: Lawrence Erlbaum Associates.

Levin, Beth. 1993. *English verb classes and alternations: A preliminary investigation*. Chicago, IL: University of Chicago Press.

Lieber, Rochelle. 2004. *Morphology and lexical semantics*. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9780511486296

Lieber, Rochelle. 2007. The category of roots and the roots of categories: What we learn from selection in derivation. *Morphology* 16(2). 247–272. DOI: https://doi.org/10.1007/s11525-006-9106-2

Lieber, Rochelle. 2010. On the lexical semantics of compounds: Nonaffixal (de)verbal compounds. In Sergio Scalise & Irene Vogel (eds.), *Cross-disciplinary issues in compounding*, 127–144. Amsterdam/Philadelphia: John Benjamins.

Lieber, Rochelle. 2015. The semantics of transposition. *Morphology* 25. 353–369. DOI: https://doi.org/10.1007/s11525-015-9261-4

Lieber, Rochelle. 2016. *English nouns: The ecology of nominalization*. Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9781316676288

Löbner, Sebastian. 1985. Definites. *Journal of Semantics* 4(4). 279–326. DOI: https://doi.org/10.1093/jos/4.4.279

Löbner, Sebastian. 2011. Concept types and determination. *Journal of Semantics* 28. 279–333. DOI: https://doi.org/10.1093/jos/ffq022

Löbner, Sebastian. 2013. *Understanding semantics*. New York, London: Routledge 2nd edn.

Löbner, Sebastian. 2014. Evidence for frames from human language. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: Applications in language, cognition, and philosophy*, 23–67. Dordrecht: Springer. DOI: https://doi.org/10.1007/978-3-319-01541-5_2

Löbner, Sebastian. 2015. Functional concepts and frames. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Meaning, frames, and conceptual representation*, 35–62. Düsseldorf: Düsseldorf University Press.

Marchand, Hans. 1969. *The categories and types of present-day English word-formation*. München: Beck 2nd edn.

Osswald, Rainer & Robert D. Jr. Van Valin. 2014. Framenet, frame structure, and the syntax-semantics interface. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: Applications in language and philosophy*, 125–156. Dordrecht: Springer.

Petersen, Wiebke. 2007. Representation of concepts as frames. In Jurgis Skilters (ed.), *The Baltic international yearbook of cognition, logic and communication* 2. 151–170. Manhattan, KS: New Prairie Press.

Petersen, Wiebke & Tanja Osswald. 2014. Concept composition in frames: Focusing on genitive constructions. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: Applications in language, cognition, and philosophy*, 243–266. Dordrecht: Springer.

Petersen, Wiebke & Thomas Gamerschlag. 2014. Why chocolate eggs can taste old but not oval: A frame-theoretic analysis of inferential evidentials. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types: Applications in language,cognition, and philosophy*, 199–218. Dordrecht: Springer. DOI: https://doi.org/10.1007/978-3-319-01541-5_9

Plag, Ingo. 2004. Syntactic category information and the semantics of derivational morphological rules. *Folia Linguistica* 38(3–4). 193–225. DOI: https://doi.org/10.1515/flin.2004.38.3-4.193

Pollard, Carl & Ivan A. Sag. 1994. *Head-driven phrase structure grammar*. Chicago, IL: University of Chicago Press.

Pustejovsky, James. 1995. *The generative lexicon*. Cambridge, MA: MIT Press.

Pustejovsky, James. 2011. Coercion in a general theory of argument selection. *Linguistics* 49(6). 1401–1431. DOI: https://doi.org/10.1515/ling.2011.039

Riehemann, Susan. 1998. Type-based derivational morphology. *Journal of Comparative Germanic Linguistics* 2. 49–77. DOI: https://doi.org/10.1023/A:1009746617055

Sag, Ivan A. 2012. Sign-based construction grammar: An informal synopsis. In Hans C. Boas & Ivan A. Sag (eds.), *Sign-based construction grammar*, 69–202. Stanford, CA: CSLI Publications.

Schulzek, Daniel. 2014. A frame approach to metonymical processes in some common types of german word formation. In Thomas Gamerschlag, Doris Gerland, Rainer Osswald & Wiebke Petersen (eds.), *Frames and concept types*, 221–242. Dordrecht: Springer. DOI: https://doi.org/10.1007/978-3-319-01541-5_10

Williams, Edwin. 1981. Argument structure and morphology. *Linguistic Review* 1. 81–114. DOI: https://doi.org/10.1515/tlir.1981.1.1.81

Wunderlich, Dieter. 2012. Operations on argument structure. In Claudia Maienborn, Klaus von Heusinger & Paul Portner (eds.), *Semantics: An international handbook of natural language meaning* 3. 2224–2259. Berlin: De Gruyter Mouton. DOI: https://doi.org/10.1515/9783110253382.2224

Zimmer, Karl E. 1964. Affixal negation in English and other languages: An investigation of restricted productivity. *Word* 20(5). 21–45.

How to cite this article: Andreou, Marios. 2017. Stereotype negation in Frame Semantics. *Glossa: a journal of general linguistics* 2(1): 79.1–30, DOI: https://doi.org/10.5334/gjgl.293

Submitted: 08 November 2016 **Accepted:** 27 Jun 2017 **Published:** 06 September 2017

]u[*Glossa: a journal of general linguistics* is a peer-reviewed open access journal published by Ubiquity Press.

OPEN ACCESS