

The treatment of emotion vocabulary in FrameNet: Past, present and future developments¹

Josef Ruppenhofer

Abstract

Both for psychology and linguistics, emotion concepts are a continuing challenge for analysis in several respects. In this contribution, we take up the language of emotion as an object of study from several angles. First, we consider how frame semantic analyses of this domain by the FrameNet project have been developing over time, due to theory-internal as well as application-oriented goals, towards ever more fine-grained distinctions and greater within-frame consistency. Second, we compare how FrameNet's linguistically oriented analysis of lexical items in the emotion domain compares to the analysis by domain experts of the experiences that give rise (directly or indirectly) to the lexical items. And finally, we consider to what extent frame semantic analysis can capture phenomena such as connotation and inference about attitudes, which are important in the field of sentiment analysis and opinion mining, even if they do not involve the direct evocation of emotion.

1 Introduction

Emotions are a core part of human experience and also well represented in language. Both for psychology and linguistics, emotion concepts are a continuing challenge for analysis in several respects. In this contribution, we take up the language of emotion as an object of study from several angles.

¹ The author was partially supported by the German Research Foundation (DFG) under grant RU 1873/2-1.

The first goal is to illustrate how FrameNet,² the project implementing an analysis of the English vocabulary in terms of Frame Semantics, ‘works’. Using emotion vocabulary as our case study, we consider how FrameNet’s analyses have been developing over time, due to theory-internal as well as application-oriented goals, towards ever more fine-grained distinctions and greater within-frame consistency.

A second goal of this study is to explore how FrameNet’s linguistically oriented analysis of lexical items in the emotion domain compares to the analysis by domain experts of the experiences that give rise (directly or indirectly) to the lexical items. Despite the skepticism expressed towards lexical analysis by some domain experts such as Scherer (2000), we want to examine to what extent the notions FrameNet uses for its analysis do match ones found in psychological theories.

Finally, we consider whether and how frame semantic analysis can capture phenomena such as connotation and inference about attitudes, which are important in the field of sentiment analysis and opinion mining, even if they do not involve the direct evocation of emotion.

The remainder of this paper is structured as follows. We first introduce Frame semantics and FrameNet in section 2. In section 3 we consider how FrameNet dealt with the vocabulary of emotion in its earliest phases. This is followed by an overview of psychological theories of emotion in section 4, with special focus on lexical theories of emotion. In section 5, we return to FrameNet and discuss newer developments in how it deals with the vocabulary of emotion. We then broaden our view beyond the core emotion vocabulary and consider issues of connotation and attitude inference in section 6. We offer a summary and conclusions in section 7.

2 Frame semantics and FrameNet

The FrameNet project represents the attempt to implement frame semantics as conceived by Charles Fillmore (1982; 1985). The theory of frame semantics is both temporally and conceptually prior to its specific embodiment by FrameNet. As a theory of linguistic semantics, it focuses on the knowledge that speakers make

² The project’s homepage is found at: framenet.icsi.berkeley.edu.

use of in producing and understanding utterances. It emphasizes continuity between language and experience. Frames are schematic representations of situations as well as their participants, props, and parameters. Words are said to evoke such frames. Lexical semantics in the frame semantic mold seeks to capture the motivation within a speech community for entertaining the relevant frame concept.

The electronic database that the FrameNet project builds contains descriptions of the semantic frames underlying the meanings of the English words being analyzed, and information on the valence (semantic and syntactic) of simple words and multi-words, derived from a sample of annotated attestations from corpora of contemporary British (BNC) and American English (ANC). The FrameNet database is intended to serve both theoretical linguistic research and applications of natural language processing (Baker et al. 1998). It is made accessible in a structured format that supports both human study and browsing, as well as machine readability.

The types of situations that FrameNet analyzes include all aspectual classes, covering states, processes, activities and relations. Some frames also focus on entities/things (e. g. *Clothing*, *Hair configuration*). The participants, props and roles of a frame can include agents, inanimate objects, elements of the setting, and properties/parameters of the situation. The syntactic dependents (broadly construed) of a predicating word correspond to the frame elements (FEs) of the frame (or frames) associated with that word.

Despite Fillmore's early work on case grammar (Fillmore 1968), FrameNet does *not* assume a set of universal semantic roles that applies to all predicates. Each FE is defined relative to a *single* frame. However, FEs in one frame can be related to FEs in another frame via frame-to-frame relations (see below), thereby allowing for the capturing of generalizations. But these connections between FEs of different frames need to be stated explicitly: identical FE names in different frames are not meant to imply identity or equivalence. Finally, the vast majority of FEs is assigned a semantic type that, broadly speaking, reflects the selectional restrictions on role fillers.

FrameNet has a net-like structure because of the relations connecting frames to each other. FrameNet employs several types of frame relations, the most im-

portant of which are INHERITANCE, SUBFRAME, PERSPECTIVE ON, and USING. INHERITANCE is FrameNet's label for the *is-a* relation. The SUBFRAME relation connects a child frame which is a subevent of a complex event to that complex parent frame. For instance, the *Criminal process* frame has subframes of *Arrest*, *Arraignment*, *Trial*, and *Sentencing*. A frame that stands in a PERSPECTIVE ON relation to another provides a particular view on an 'unperspectivized' parent that is an alternative to at least one other kind of view. The *Hiring* and *Get a job* frames, for instance, take the employer's and employee's view, respectively, on their parent frame *Employment start*. In the USING relation, the child frame presupposes the parent frame as background. The *Volubility* frame for instance uses the *Communication* frame.

Frame relations are accompanied by parallel, but potentially partial, mappings between the frame elements of the frames that are in a relation. These FE mappings allow one to, for instance, recognize the identity between a *SUSPECT* in the *Arrest* frame and a *DEFENDANT* in the *Arraignment* frame.

A Lexical Unit (LU) is the pairing of a morphological lemma with a meaning; a word sense. The meaning of a LU is partially expressed by the relation between the lemma and a FN frame, i.e. between lexical form(s) and the semantic frame they evoke. FrameNet covers LUs of any lexical category (verbs, nouns, adjectives, prepositions, etc.). The *Similarity* frame, for instance, includes among its LUs the preposition *like*, the adjective *similar*, the verb *resemble*, and the noun *difference*. Multi-word expressions (MWEs) are also possible LUs for FrameNet to include: the *Similarity* frame includes the MWE *spitting image*. As can be seen from the list of lexical units in the *Similarity* frame, FrameNet puts antonyms in the same frame (e. g. *similar*, *different*). The inclusion of *spitting image* illustrates that the members of a frame are not all equally general or specific: *spitting image* applies to similarity of appearance only, unlike the other LUs.

The semantic analyses of FrameNet are based on data and documented through annotations on example sentences. Such annotations are performed in one of two modes. In the *lexicographic* mode, FrameNet annotates instances of particular LUs from a corpus. The instances to be annotated are sampled from the corpus in such a way that the full set of combinatorial possibilities (in the relevant corpus) is captured. The annotations are, however, not meant to be statistically representative: frequent patterns of FE realization do not receive more annotations than less

frequent ones. These lexicographic annotations do not fully cover all frame-evoking words in the sentences, let alone documents, in the underlying corpora (BNC, ANC). However, FrameNet has another mode of annotation, *full-text*, in which every instance of every frame-evoking word and its FEs are labeled on specific documents, resulting in dense annotations that show how frames embed each other and collocate.

Annotated instances of LUs record the following information about the LUs' syntactic dependents: their FE role, their phrase type, and their grammatical function. By way of example, let us briefly consider FrameNet's *Revenge* frame. Its definition mentions the core frame elements (which bear frame-specific, non-universalist names) and their relationships within the frame:

This frame concerns the infliction of punishment in return for a wrong suffered. An **Avenger** performs a **Punishment** on an **Offender** as a consequence of an earlier action by the **Offender**, the **Injury**. The **Avenger** inflicting the **Punishment** need not be the same as the **Injured Party** who suffered the **Injury**, but the **Avenger** does have to share the judgment that the **Offender's** action was wrong. The judgment that the **Offender** had inflicted an **Injury** is made without regard to the law.

Each FE is also defined and exemplified separately, though we do not show this here so as to conserve space. In the FrameNet database, the *Revenge* frame is recorded to be a subtype (by INHERITANCE) of the frame *Rewards and Punishments*. Its siblings include the frame for *Fining*, which unlike the *Revenge* frame is defined against a legal background and involves punishments in the form of financial levies. Some sample annotations in the *Revenge* frame are:

- (1) One day [she **AVENGER**] would **get even** [with Donna **OFFENDER**] [for this **INJURY**]. [INI **PUNISHMENT**]
- (2) One day [I **AVENGER**]’ll **get back** [at you **OFFENDER**]. [DNI **INJURY**] [INI **PUNISHMENT**]
- (3) The next day, [the Roman forces **AVENGER**] took **revenge** [on their enemies **OFFENDER**]. [INI **PUNISHMENT**] [DNI **INJURY**]
- (4) [The ban **PUNISHMENT**] is [Prince Charles’s **AVENGER**] **revenge** [for her refusal to spend Christmas with the rest of the royals **INJURY**]. [DNI **OFFENDER**]

In examples (1)–(4), the targets are marked in bold font, support verbs appear underlined, and FEs are enclosed in square brackets. The group of examples includes verbal and nominal targets. It also illustrates that the realization of FEs may vary quite a lot within a frame. For instance, the FE **OFFENDER** is marked by three different prepositions in examples (1)–(3). As these last two observations

indicate, FrameNet's analyses provide resources for recognizing paraphrases. This capability has, for instance, been exploited in the automatic recognition of textual entailment (Burchardt 2008).

Further, since frames are understood to be cognitive representations of situations, rather than just specific linguistic meaning representations, they are assumed to possess considerable generality across cultures and languages. While this generality is known to have its limits (e. g. Boas 2009), it is nonetheless broad enough to have 1) spawned sister projects to FrameNet for several other languages (among them, for instance, Spanish, Swedish, Japanese, Brazilian Portuguese) and 2) given rise to considerable interest in Frame semantics as an interlingua for translation. Speaking in an oversimplified manner, one might say that aligned FrameNet resources for multiple languages allow one to approach translation as a problem of cross-language paraphrase.

It is important to keep in mind that historically the main focus of FrameNet has been on frames that can be lexically evoked, as valence is its key concern. Accordingly, while FrameNet does include some so-called non-lexical frames which cannot be directly evoked by lexical material but which are relevant for capturing background relevant to related groups of frames, that state of affairs is the exception rather than the rule. FrameNet does therefore not deal with scripts in the sense of Schank and Abelson (1977) such as the restaurant script, for which no sizable amount of specific lexical material is available. Finally, note that FrameNet is for the most part concerned with general language vocabulary. Some work has, however, been done on specific domains and more efforts along these lines exist outside of FrameNet (e. g. BioFrameNet (Dolbey et al. 2006), Kicktionary (Schmidt 2009)).

3 Original treatment of emotion vocabulary in FrameNet

Emotion vocabulary is a core and, in terms of quantity (Hobbs and Gordon 2008), well-developed part of a language's vocabulary. Not surprisingly, emotion was

one of the 13 general domains that the FrameNet project set out to handle when it began.³

In the first two releases of the FrameNet database (1.2 and 1.3), the emotion vocabulary was grouped into frames by several criteria being applied separately and in parallel. Some frames were, unsurprisingly, motivated by valence patterns (sentence perspective). For instance, the frame `Experiencer subj` was distinguished from `Experiencer obj` since for the former group of predicates the experiencer is realized as a subject (5) and for the latter it is realized as an object (6) in active-form sentences.

- (5) [I `EXPERIENCER`] **like** ice-cream.
- (6) Ice-cream **pleases** [me `EXPERIENCER`].

Other frames, picked up on salient metaphorical conceptualization. For instance, the frame `Emotion heat` held lexical units such as *boil.v*, *simmer.v*, *stew.v*, which are all motivated by the EMOTION IS HEAT metaphor (Lakoff and Johnson 1980). Still other frames were motivated by particular emotion concepts: for instance, the frame `Desiring` contains lexical units expressing forms of desiring.

- `Experiencer subj` (love.v, like.v, hate.v; fear.v; resent)
- `Experiencer obj` (please.v, displeasure.v; anger.v)
- `Emotion directed` (*angry.a at*)
- `Subject stimulus` (satisfying.a; pitiful.a)
- `Desiring` (hope.v, desire.v, desire.n, covet.v, hankering.n)
- `Judgment` [Judgment direct address, Judgment communication] (*reproach.v* : *approve.v*)
- `Desirability` (okay.a; nasty.a, pitiful.a; great.a, marvellous.a)
- `Emotion heat` (stew.v, simmer.v, boil.v; burn.v)

It is clear that in some cases the original frame organization assumes a relatively unspecific shared background given the considerable differences between the LUs included in particular emotion frames. For instance, the `Experiencer obj` frame contained the verbs *please* and *anger*, which reference rather different kinds of emotions. Likewise, and maybe more surprisingly, the adjective *angry*

³ The full list of domains includes HEALTH CARE, CHANCE, PERCEPTION, COMMUNICATION, TRANSACTION, TIME, SPACE, BODY (parts and functions of the body), MOTION, LIFE STAGES, SOCIAL CONTEXT, EMOTION and COGNITION (Baker et al. 1998).

and the noun *anger* belong to a different frame from the verb *anger*, even though they all reference the same emotion.

Observations such as these made it desirable to consider a reorganization of how the emotion vocabulary is structured into FrameNet frames. However, before we consider what developments have been ongoing in FrameNet’s modeling of the emotion domain, we will go on an excursion into psychological theories of emotion. While FrameNet clearly seeks to capture laymen’s (folk) understandings of their experiences, it is instructive to consider how domain experts conceive of emotions, especially because it has always been controversial how a psychological theory of emotion should relate to a linguistic analysis of the emotion vocabulary. Notably, while speakers may talk about their emotions in what may be (mere) folk categories, their linguistic productions constitute the bulk of the primary data that psychologists studying emotion have available.

4 Psychological theories of emotion

In this section, we briefly survey psychological theories of emotions at a high-level of abstraction, following the discussion in Scherer (2000), and then look more closely at one particular subtype, lexical models.




Figure 1: The circumplex model (= Figure 2 of Russell 1980,1167)

Dimensional approaches propose that one or more dimensions such as valence or activation/arousal are sufficient to capture the most important differences between emotions. A well-known instance of multi-dimensional models is the circumplex model of affect (e. g. Russell 1980; Plutchik 1980), in which emotions are arranged on two dimensions, valence and arousal (cf. Figure 1). The model is intended to capture several key aspects of emotions (or emotion terms).⁴ First, many emotions are associated with oppositions (*happy-sad*). Second, most of the variance in rating studies where similarity of emotions (emotion terms) is judged can be accounted for by the two dimensions of valence and arousal. Finally, the different emotions (emotion terms) can be differentiated and do not simply lie on top of each other along the axes in the space of emotions.

Categorical approaches (also: discrete emotion approaches) assume identifiable distinct types of emotions. A well-known instance of this approach is Ekman (1992), who suggests that there are six basic emotions (happiness, sadness, anger, disgust, fear, surprise) that people across cultures can recognize in others based on their facial expression. The basic set of emotions is assumed to have proven useful as evolutionary adaptations. Besides the basic emotions, additional emotion types are believed to arise as combinations of the basic types. As suggested by Scherer, one piece of evidence favoring discrete emotion-approaches is the fact that high-frequency verbal labels exist that seem to refer to prototypical emotions (anger, sadness, etc.).

Meaning-oriented approaches constitute a third type of psychological model. These rely on the structure of the semantic field of a language's emotion terms to build a psychological model. One subtype of meaning-oriented model is the lexical model. While Scherer notes the intuitive appeal of this approach as it accesses what he calls "common cultural interpretation patterns" (148), he casts doubt on such approaches ultimately leading to a better understanding of emotions:

The linguistic labels attached to specific types of affective states are not always helpful. As is true for many other areas of psychology, popular usage of some terms

⁴ Note that as Figure 1 shows, in Russell's work the stimuli were adjectives. They are treated as "emotion names". There is no explanation how the emotions evoked by verbs like *love*, *please*, etc. might fit into the analysis.

has created semantic constructs that are less than optimal for exact scientific description. (Scherer 2000, 142)

In addition, according to Scherer, there is a social-constructivist strand of emotion research, which assumes that the meaning of emotion depends on socio-culturally constructed patterns of behavior and values. A language's lexicon is taken to reflect the speakers' social constructions of emotional meaning.

Componential approaches proceed from the joint assumptions that emotions result from a cognitive evaluation of a situation (though not necessarily a conscious evaluation) and that this evaluation produces a pattern of reactions in different domains of response (physiology, action tendencies, feeling). In his own componential approach, Scherer defines emotions as follows:

[...] emotions are episodes of coordinated changes in several components (including at least neurophysiological activation, motor expression, and subjective feeling but possibly also action tendencies and cognitive processes) in response to external or internal events of major significance to the organism. (Scherer 2000, 139)


Scherer points out that different theories within the componential approach differ on how differences between emotions are conceived of and how many different emotions there are. On the relation between emotions and emotion labels in language, he says:

The bases of verbal labels of emotional states are the changes in conscious subjective feeling states. Although the feeling states may reflect all the changes characterizing an emotion process in all of the organismic subsystems, verbal labels often represent only a salient part of those changes, those that reach awareness. (Scherer 2000, 150)

Given our linguistic interests, we will now look more closely at some lexical models, Scherer's skepticism of their value for psychological theory-building notwithstanding. For a linguistic theory of the emotional vocabulary, the psychological theories that looked most closely at linguistic evidence for emotional experience are clearly relevant. And given the frame semantic perspective we are taking, the "common cultural interpretation patterns" that are activated by emotion terms are what we are after (Scherer 2000, 148).

4.1 Ortony et al.'s models

We start by considering Ortony et al.'s (1987) analysis of what they call the affective lexicon.



The authors propose the taxonomy displayed in Figure 2. Note that the predicates that they analyze are overwhelmingly adjectives, deverbal participles included among them. Verbs, in particular non-stative verbs, are underrepresented.

Figure 2: Ortony et al.'s taxonomy of affective conditions (= Figure 1 of Ortony et al. 1987, 349)

In their analysis, the most prototypical examples of emotion terms are those that (a) refer to internal, mental conditions – in distinction to external and physical ones; (b) are states as opposed to actions or properties/dispositions; and (c) which focus on affect rather than on behavior or cognitions as their referential focus. Importantly, Ortony et al. (1987) make a distinction between words that *refer* to

emotions and words that *implicate* emotions.⁵ Ortony et al. propose a test for what constitutes a genuine emotional term. The adjective *happy*, for instance, is considered to be a genuine emotion term because subjects rate both “feeling happy” and “being happy” as emotions. By contrast, the term *ignored* is not taken to be a genuine emotional term, because while subjects consider “feeling ignored” an emotion, they do not do so with respect to “being ignored”.

In subsequent work, Ortony et al. (1988) lay out a model for the cognitive structure of emotions. This model is intended to be one of the emotions, rather than of the lexical items people use to refer to emotions (as they understand and experience them). On the other hand, for this model of emotions, language is used as evidence for the existence of particular emotions. Note that the typology of Ortony et al. (1988) has influenced research in Artificial Intelligence (Elliott 1992, Hobbs and Gordon 2008) that modifies and extends the model. However, we will not consider this subsequent work here.

One goal of Ortony et al.’s model is to capture the qualitative differences between the emotions that dimensional models typically obscure. The authors also want to explain how people’s construals of their experiences cause them to experience certain emotions. Further, Ortony et al. are interested in how different emotions are related to each other, and what the structure of the individual emotions is. Accordingly, they focus on these particular aspects:

- eliciting conditions,
- variables that influence the intensity of emotions,
- distinguishing affective and mental conditions that are genuinely emotional from those that are not.

The organization of emotions in Ortony et al.’s model is shown in Figure 3. At the highest level, the emotions are grouped into three broad classes, reflecting their focus on one of three salient aspects of the world that the authors assume are central: *events*, *agents*, and *objects*. Event-based emotions are said to be based on goals. Agent-based emotions are taken to reflect the fact that we attribute re-

⁵ This is the central concern of these authors. The taxonomy they develop is seen as more of a by-product of their discussion of what does and does not refer to an emotion. Another important concern for Ortony et al. (1987) is the question whether emotions are prototype concepts or not, an issue we will not take up here.

sponsibility to agents in terms of standards of behavior. And object-based emotions are taken to be based on attraction, which is taken to ultimately be rooted in attitudes.

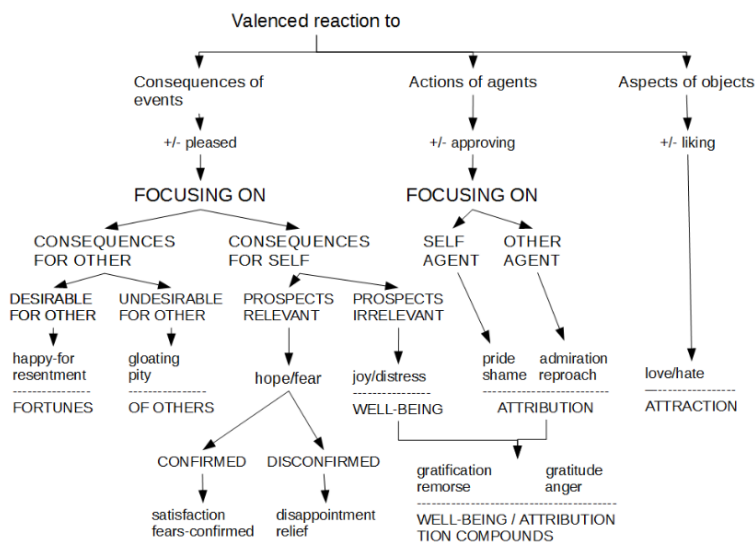


Figure 3: Emotion typology of Ortony et al. (1988, 19)

Note that in Ortony et al.'s model what is referred to as intensity is the actual intensity of instances of emotional experience rather than what one might call lexical intensity. The fact that *love.v* communicates a more intense emotion than *like.v* is not what they focus on. Rather, different instances of loving can be experienced with different intensity depending on various global and local factors. Global factors are ones that play a role for any kind of emotion, such as the experiencer's level of arousal. Local factors are ones that come into play only for particular subtypes of emotions. For instance, the degree to which one is *happy-for* somebody else depends (among other things) on the degree to which one likes that person: I might be happier if my wife or my brother wins the lottery than if my neighbor or colleague wins.

4.2 Appraisal theory

As a second lexical model, we will consider Appraisal theory (Martin and White 2005), which is a theory within Systemic Functional Linguistics (SFL) rather than psychology. We include it in this discussion because it represents a sizable body of work and because it has been influential on sentiment analysis.

From the point of view of SFL, language is a resource for relating ideational, interpersonal and textual meaning to each other. Appraisal is understood to be about interpersonal meanings concerned with negotiating social relations. Appraisal theory consists of three sub-theories for affect, appreciation, and judgment, as shown in Figure 4. We will discuss these sub-systems in turn.

4.2.1 Affect

Affect is a basic “system” within the theory. It covers the expression of emotions of any kind (happiness, fear, etc.). According to Martin and White (2005, 42) “[a]ffect is concerned with registering positive and negative feelings: do we feel happy or sad, confident or anxious, interested or bored?”

Affect is subdivided further by various criteria. One criterion is positive/negative valence. A second is whether an expression refers to a para-/extralinguistic manifestation (*sob, weep*) vs. an internal experience (*upset*). A third criterion differentiates between triggered/reactive experiences vs. moods and general dispositions. The intensity of feeling constitutes a fourth factor. Another is the question whether the experience is due to a “realis” context (e. g. a reaction such as *dislike.v*) vs. an irrealis context (e. g. an intention or an attitude towards a possible event such as *fear.v*). As a sixth criterion, the theory distinguishes 3 major affect groups:

- un/happiness (“affairs of the heart”)
- dis/satisfaction (emotions related to the pursuit of goals)
- in/security (“ecosocial well-being”)

4.2.2 Appreciation

Appreciation focuses on things; it is taken to concern aesthetics. It is said to be about feelings that are institutionalized as propositions which concern a) reactions to things (‘do they catch our attention?’, ‘do they please us?’), b) their composition (‘balance and complexity’) and c) their value (‘how innovative, timely, authentic ...’).

The linguistic prototype that expresses appreciation is as follows: “APPRECIATION construes attitudes about texts, performances and natural phenomena, and fits into frames such as *I consider it x*: ... I consider it *innovative/unimaginative*.” (Martin 2003, 173).




Figure 4: Overview of appraisal resources (Martin and White 2005, 38)

4.2.3 Judgment

Judgment concerns attitudes about people and their behavior; it is thus related to questions of ethics. It is concerned with positive and negative judgments of behavior / actions. Judgment is said to be about feelings “institutionalized as proposals”. Judgment is taken to have two dimensions. First, in what social esteem is somebody held based on how normal, capable, and/or dependable they are? Second, what kind of social sanction is placed on somebody with respect to their decency and truthfulness?

The linguistic prototype for Judgment is the following: “JUDGMENT construes attitudes about character, designed to sanction or proscribe behavior, canonically in the grammatical frame *It was ‘x’ of/for her/him to do that*” (Martin 2003, 173).

4.3 Summary

The different psychological and linguistic analyses that we have considered highlight various aspects that could be taken into account in a full analysis of emotional vocabulary. A first aspect to consider is the relation of an emotional term to terms referencing other psychological concepts, in particular the distinction to cognition; alternatively, the focus on affect vs. behavior vs. cognition. This aspect is an important part of Ortony et al.'s 1987 work. Another aspect is the question whether there are basic emotions or emotional primes, on which the analysis of other, more complex emotions should build. The analysis of emotional vocabulary proposed by Johnson-Laird and Oatley (1989) posits five basic emotions for English as primitive notions without further internal semantics. (They roughly correspond to five of Ekman's basic emotions.) Wierzbicka's Natural Semantic Metalanguage framework is another well-known semantic theory that assumes semantic primes. However, there are no specific emotional primes present in her inventory other than *feel* (Wierzbicka 1992, 541). By contrast, neither Ortony et al.'s models nor Appraisal Theory assumes primitive emotions. Thirdly, an analysis might model a typology of eliciting conditions that are assumed. Ortony et al.'s 1988 model puts forth a well-developed proposal of this sort. Fourthly, the intensity of emotional experience and/or the intensity of lexical expressions could be specified as part of the analysis. Appraisal theory pays attention to (lexical) intensity. For instance, Martin and White (2005, 48) assign low, medium and high intensity to the verbs *dislike*, *hate* and *detest*, respectively. Intensity also is a key aspect of Ortony et al.'s 1988 model. However, they are more concerned with contextualized intensity of experiences that results from both local and global influences. While these influences are relevant to text understanding and may be signaled linguistically, they are independent of the inherent lexical semantics of the emotion terms. A fifth point to consider is the difference between personality traits and phases / episodes of emotional experience. This distinction is present in Appraisal theory (also in Johnson-Laird and Oatley's (1989) model), whereas the theories of Ortony et al. mostly seem to focus on the phases, setting aside the analysis of dispositions. And lastly, the scope taken by lexical items on the overall scenario from eliciting event to the set of coordinated resulting changes of emotional experience should be represented. For instance, does a term refer to a resulting state (e. g. *sad.a*) or does it focus on the transition into the resulting state

(e. g. *sadden.v*)? This distinction seems to significantly overlap with Appraisal theory's distinction between what it calls mental processes (such as *Pat hates Kim*) and relational states (*Pat is unhappy*).

In what follows, we consider how these considerations are, or could be, addressed by the frame semantic analysis of emotion vocabulary by the FrameNet project.

5 Emotion vocabulary in FrameNet

If we consider the FrameNet database at the time of this writing (May 2015), it is clear that the frame organization implemented focuses on different criteria than, for instance, the taxonomy of Ortony et al. 1988 (cf. Figure 3). For instance, the frame-mate verbs *please* and *anger* in the `EXPERIENCER obj` frame are in different parts of Ortony et al.'s 1988 taxonomy. Likewise, as discussed previously, the adjective *angry* and the verb *anger* belonged to two different frames, although the emotion they reference is the same. Conversely, making these comparisons shows that the taxonomy does not take into account the valence patterns associated with the predicates. In fact, the analyses of Ortony and colleagues (1987, 1988) purposely represents many emotion terms only indirectly through one of a set of related terms. E. g. the 1987 taxonomy favors stative predicates and thus the adjective *angry* is part of it but not the verb or the noun *anger*.

However, we cannot look at the current organization of emotion vocabulary in FrameNet as reflecting a single synchronous analysis. Rather, the frames related to emotion reflect different phases in FrameNet's development. As an ongoing and living project, FrameNet has over time developed new (and more explicit) criteria as to how frames should be delimited and distinguished from each other (Ruppenhofer et al. 2010). Overall, these new developments lead to smaller, internally more consistent frames than were constructed in the earliest phases of the project. An important goal was that the frame organization would be more suitable for use in natural language processing systems dealing with tasks such as textual entailment recognition, paraphrasing, translation, etc. The core criteria that were adopted are the following:

- The LUs in a frame should take the same number and types of syntactic dependents. *Scary.a* and *fear.v* can thus not share a frame as the former does not obligatorily modify, or predicate of, an EXPERIENCER.
- The LUs should belong to the same Aktionsart / pick out the same phase of complex events.
- The LUs should have the same entailments. Accordingly, *shoot.v* is a member of the Hit target frame while *kill.v* and *behead.v* are in the frame Killing.
- The LUs should take the same participant's perspective on the event (in active-form sentences). Applying this criterion, *buy.v* *sell.v* are in different frames, Commerce buy and Commerce sell, whereas in earlier frame-semantic theorizing (Fillmore 1976) they shared a frame.⁶
- The FEs should have the same relations to each other across all lexical units. This criterion also argues for separating *buy.v* and *sell.v*: the PURPOSE of the BUYER is not the same as the PURPOSE of the SELLER.
- The presuppositions and future-oriented expectations associated with the lexical units should be the same. Accordingly, *examine.v* and *cross-examine.v* cannot share a frame as the latter evokes the context of a court proceeding, unlike the former.

In addition to positive criteria for frame distinctions, FrameNet also established which criteria it would not use to divide up frames.

- Frames should be organized to maximize within-frame suitability for paraphrasing. This criterion argues for frames including mostly (near) synonyms. It calls for the splitting up of such large frames as EXPERIENCER subj and EXPERIENCER obj that are based on argument linking.
- Syntactic variants that arise through the combination of LUs with productive constructions (e. g. the passive or the middle) should stay within the same frame.
- Differences in intensity do not call for frame divisions: *good.a* and *great.a* both belong to the frame Desirability.
- Differences due to deixis, register, dialect/variety should not result in frame differences. Thus, *botch.v*, *mess up.v*, *fuck up.v* are all members of the Bungling frame.

⁶ If one takes into account FrameNet's hierarchy, then it is clear that both LUs still evoke the Commercial_transaction frame that connects all related sub-frames.

- Differences in valence do not lead to different frame membership. Accordingly, *criticize.v* and *praise.v* are both members of the `Judgment communication` frame.⁷

If we look at the more recent analysis of emotions in FN release 1.5, we see that as a consequence of the above desiderata already there are finer divisions among the emotion frames. As an example, consider the child frames of the `Emotions by stimulus` frame, shown in Figure 5.

Some of the frames that appear in this analysis can be mapped onto groups that are found in Ortony et al.'s analysis. For instance, the frame `Other situation as stimulus` captures the idea of "Consequences for other". Similarly, the frame `Emotions success or failure` captures much of what is meant by "Prospects relevant (Confirmed)". However, as the data release of FN 1.5 represents only a snapshot of an ongoing project, it should come as no surprise that the application of the criteria is not fully realized in the domain of emotion-vocabulary. For instance, the frame `Experiencer obj` still persists, even though includes a very heterogeneous set of predicates with respect to eliciting conditions and (emotional) valence.

In the next cycle of evolution, FrameNet will see a reanalysis of the emotion area that aims for the design below, illustrated using the examples of anger- and depression-related lexical units.

Use of unperspectivized frames, with frames for particular profilings:

- event focus
 - State: *(feel) anger.n, have one's undies in a bundle; depression, have (got) the blues, be in a funk*
 - Inchoative: *become enraged, get angry, become depressed*
 - Causative: *anger.v, tick off.v, piss.off; depress.v, get sb down.v*
- experiencer focus, with distinctions for
 - Phase: *angry.a, enraged.a; depressed.a*
 - Disposition/Mood: *irascible.a; depressive.a*
- stimulus focus: *enraging.a; depressing.a*

This re-organization of the emotion vocabulary captures many of the dimensions by which emotions can be organized. The *eliciting condition* and their construal

⁷ But note that adhering to this criterion necessarily leads to inconsistencies with respect to paraphrasability.

by the experiencer are packaged in the unperspectivized frame. For instance, anger is an emotion that arises when the experiencer perceives a wrong that calls for a response. Fear is an emotion that arises when danger is perceived. Beyond the eliciting conditions, the particular scope on the overall scenario is also consistent per frame. For instance, causatives are distinguished from states. Note, that the eliciting situations are themselves frames and as such could be related to the relevant emotion frames. For instance, the *Fear* frame might be related to the *Danger* frame via the *USE* relation. Furthermore, emotion frames can be referred to by frames that encode actions that experiencers are likely to engage in. Accordingly, the *Fleeing* frame can *USE* the *Fear* frame. Through these links, the idea of *action tendencies* in the psychological literature can be captured. The frames also capture the difference between personality traits and episodes of emotion, which previous models do not deal with.⁸ The valence of emotions is also captured, albeit only implicitly, through membership in the (leaf)-frames.




Figure 5: Frames inheriting from Emotions by stimulus

Basic lexical intensity (e. g. *love.v* > *like.v*) is not represented (yet). Also, FrameNet does not capture the knowledge about which local variables (in the sense of Ortony et al. 1988) influence the intensity of instances of particular subtypes of emotional experience. As noted above, Ortony et al. (1988) suggest that the degree to which one is *happy-for* somebody else depends (in part) on the degree to which one likes that person.

Also unaddressed by the frame semantic analysis is the idea of qualia or the subjective feeling an experiencer has when experiencing a particular emotion.

⁸ Of course, phases of anger and depression still differ considerably in length from each other, the former being much shorter-lasting.

However, the inventory of LUs that can refer to particular emotions may capture some of the “coordinated changes” in an experiencer that make up the qualia of an emotion. For instance, *having goosebumps* or *having one’s hair stand on end* may simply be reactions to cold temperature. But they also are typical physical manifestations of particular kinds of emotions involving arousal, and they are frequently used metonymically to evoke these emotions.

Finally, if we compare FrameNet’s intended restructuring with Ortony et al.’s (1988) organization, we see that one key difference between them is that in Ortony et al.’s theory the profilings are not modeled. It captures only the higher-level groupings represented by the unperspectivized frames. Another big difference is that through its use of frames and frame-to-frame relations, the FrameNet analysis addresses not only the modeling of the emotion vocabulary but also the modeling of the vocabulary of eliciting conditions and actions that (typically) result from emotional experiences.

6 Beyond core emotion vocabulary

There are aspects of affective language that earlier linguistic and psychological theories of emotional language omit from their model. For instance, analyses of core emotion language tend to leave aside predicates that evoke scenes where the emotional state is presupposed or backgrounded. For instance, *laud* is speech expressing praise (cf. (7)). This verb and other verbs like it are missing from the list of terms analyzed by, for instance, Johnson-Laird and Oatley (1989).

(7) Bermuda’s late former Chief Justice Sir Richard Ground, who died last month at the age of 64, was yesterday **lauded** at a packed special sitting of the Supreme Court.

FrameNet’s current representation relates the `Judgment communication` frame that hosts the verb *laud* to the frame `Judgment` via the `USE` relationship. However, that connection necessarily leaves the parameter of valence/polarity out of its purview because both frames contain items that differ in terms of their polarity. For instance, both the positive *praise* and the negative *criticize* are hosted by `Judgment communication`.

A large issue is what traditionally is called connotative meaning. It is outside the purview of the lexical models of emotions discussed previously in this section.

FrameNet covers some of this information. For instance, the noun *coot* in the frame *People by age* has a Semantic type *Negative judgment*. However, the representation is inexplicit: we do not know whose judgment this is (the speaker's) and which FE it concerns (the FE PERSON). And, more generally, connotations are often not represented in a machine-readable format. For instance, for the verb *boast.v* the verbal definition, which FrameNet takes from the Concise Oxford Dictionary, indicates to human users that there is negative judgment but there is no semantic type associated with the LU that could be used by a machine. On the other hand, it is not clear what *specific* emotions, if any, terms with connotations like *coot.n* and *brag.v* should be tied to. If one calls somebody's behavior bragging, does that usually mean one does not like the person?

Interestingly, connotations and denotations of emotion are not in complementary distribution. Quite a few expressions carry more than one level of affective meaning at the same time. For instance, the adjective *infatuated* expresses a person's enthusiasm for a person or entity but typically implies that the enthusiasm is excessive or unwarranted in the view of the person who uses the word *infatuated* to report another person's feelings (8).

(8) Indeed, **infatuated** portfolio managers have begun to court stocks once again, and stock prices have climbed almost 10 per cent from their lows in early-July, as impressive second-quarter earnings reports soothed equity investors' double-dip fears.

Cases like this call for a more detailed representation than the use of the labels *Positive* and *Negative judgment* affords. If one simply added both tags to *infatuated*, it would not be possible to distinguish the layering of attitudes—one by a FE towards another FE, one by the speaker/reporter towards an FE – that applies in the case of *infatuated* from the co-presence of contradictory attitudes *within* the frame that is evoked by e. g. *ambivalent*:

(9) Over the years I've given a number of talks on Rachmaninov, about whose music I have felt **ambivalent** at various points in my life.

Proposals for explicitly adding information about attitudes between participants or by the speaker relative to the event or to its participants have been made by Maks and Vossen (2011) and, specifically for FrameNet, by Ruppenhofer and Rehbein (2012) and Ruppenhofer (2013). The latter authors propose extending Frame-

net by a SentiFrameNet add-on, which contains several additional representational constructs to support fine-grained sentiment analysis. The most important of these are *opinion frames*. These would contain the opinion roles SOURCE and TARGET, and they would always be tied to ‘content’ frames via a new type of frame relation, with attendant mappings between opinion roles and ‘regular’ FE roles. The opinion frames would also capture valence (or: polarity) and intensity. Accordingly, they would need to be LU-specific, and they would also need to link to LU-specific ‘content’ frames, which in turn would be linked to the existing ‘content’ frames by inheritance.

A further frontier is the issue of what lexical information might support (de-feasible) inferences about attitudes that are not expressed explicitly but *implied*. In recent years, Reschke and Anand (2011) and Wiebe and colleagues (Deng et al. 2013, Deng and Wiebe 2014, Wiebe and Deng 2014) have generated a lot of interest in pursuing this line of research. Consider an example given by Wiebe and Deng 2014:

(10) Obama will finally bring **skyrocketing** health care costs under control.

Arguably, *skyrocketing* is the only inherently evaluative lexical item in sentence (10). However, as Wiebe and Deng note, given knowledge about other lexical material in the sentence, most readers will infer that the author/speaker probably views Obama favorably in this context. The line of reasoning is basically as follows. Skyrocketing costs are viewed negatively by the speaker. This is explicitly coded by *skyrocketing*. Additionally, one may assume, or context may make it clear, that the speaker is a simple health care consumer with no vested interest in the financial success of health providers or insurers, so that keeping costs down must be her priority. If costs can be counteracted (brought under control) that is bad for them.⁹ Given a logic along the lines of my-enemy’s-enemy-is-my-friend, the speaker views such an event positively. Further, if a good or bad event is recognized as caused by another person, then we value that person accordingly. Since, in the example, the curbing of costs is brought about by Obama, he should be given credit for it.¹⁰

⁹ For the sake of the argument, we will ignore the possibility that “bring X under control” might lexically specify negative judgment of the speaker towards X.

¹⁰ We note here that Wiebe and Deng’s current model treats the speaker’s views on the participants as a binary variable. Based on the discussion of global variables influencing intensity in Ortony et

Two things are worth noting here. First, the reasoning schema that is used is one that also appears in Ortony et al.'s model of the emotional lexicon. The reasoning about Actions of Agents tracks closely with the reasoning needed for the inference in (10) and applying it to the situation depicted would suggest that the speaker, in so far as she feels affected by Obama's action, should feel *gratitude* towards him. (Another possibility is a feeling of admiration, to which personal benefit is less central.) In other words, the reasoning schemata that are evoked when emotions are *reported* could also be used to form inferences of potential, unexpressed emotions based on mentions of potential eliciting conditions. However, the analysis goals of Wiebe and Deng (2014) stop short of pursuing inference of specific implicit emotions. They focus on recognizing implicit valenced attitudes rather than specific mental/emotional states. Work that pursues the recognition of specific emotions includes, among others, Balahur et al. (2011) and Cambria et al. (2012). However, these latter efforts neither use the same semantic features as Wiebe and Deng's work, nor the same inferential calculi.¹¹

The second point to make about work like Wiebe and Deng's is the question whether the lexical information needed to support the inferences can be reduced to well-known linguistic features. The terminology in the field of sentiment analysis is not always clear in this respect. For instance, Wiebe and colleagues say that one needs to know which participants an event is *good-for* or *bad-for*. Here, one may wonder whether the notion of *good-for/bad-for* can be related to the linguistic notion of affectedness, on which there exists a large literature (for a summary, see Beavers 2011).

In any event, the inferencing mechanisms that Wiebe and Deng (2014) discuss rely on deep lexical semantic knowledge. Even if FrameNet does not currently provide all of it, we argue that it is uniquely suited to host and aggregate such

al.'s 1988 model, we might expect for example (10) that on a graded treatment, the more one likes Obama and/or the more one values health care, the more credit one will give to Obama.

¹¹ But note that the emotion recognition systems of Balahur et al. (2011) use semantic role labeling as one of its components. This suggests that the current FrameNet could be useful to emotion recognition systems even in its 'simple' function as a resource for semantic role labeling, without any extensions. This also applies to Wiebe and Deng's approach, which also makes use of semantic roles.

knowledge since a lot of it is related in the end to semantic roles resp. FEs. Affectedness (or *good-for/bad-for*) and its complementary notion of causation are very salient cases in point.

7 Conclusion

In this contribution, we have pursued several goals. On the one hand, we wanted to illustrate how FrameNet ‘works’, in particular, what kind of analyses it produces and has produced. Using emotion vocabulary as our hook, we have illustrated how FrameNet’s analyses have developed over time, due to theory-internal as well as application-oriented goals, towards ever more fine-grained distinctions and greater within-frame consistency in the attributes of LUs. Related to this, an important practical caveat for users of FrameNet’s data is to be aware that the database is not static. This means that analyses change and a snapshot of the database does not necessarily represent a state where all areas of the vocabulary covered have been analyzed exactly by same principles. This may be unfortunate but it is in the nature of an evolving research project.

A second goal of this contribution was to explore how in the domain of emotions FrameNet’s linguistically oriented analysis of *lexical items* compares to the analysis by domain experts of the *experiences* that give rise (directly or indirectly) to the lexical items. Despite Scherer’s (2000) skepticism towards lexical analysis, it seems that the notions FrameNet uses for its analysis do match ones found in psychological theories (in particular that of Ortony et al. 1988), though in particular the notion of *profile* that the linguistic analysis requires may not have any kind of counterpart in the psychological theories. As we suggested, a greater alignment of psychological theory and FrameNet’s analysis could be achieved if additional notions such as eliciting conditions and action tendencies were cashed out in FrameNet more often via frame relations.

Finally, we argued that when going beyond vocabulary that evokes emotional experiences directly and considering issues such as connotation or inference, the schematization of emotional experiences by FrameNet and psychological theory still proves useful. We suggested that descriptions of eliciting conditions can be used to infer potential emotional experiences, assuming, in particular, that crucial knowledge about the harm and benefit of events is provided. We believe that

FrameNet has the required knowledge or can provide it. The need of sentiment analysis for deep lexical knowledge thus makes a powerful case that FrameNet needs not only to be widened in terms of the lexical items covered but also to be deepened in certain ways. We believe that other (related) language processing tasks such as metaphor detection and understanding would benefit from such deepening as well.

References

- Baker, Collin F. / Fillmore, Charles J. / Lowe, John B. (1998): The Berkeley FrameNet project. Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics – Volume 1. Stroudsburg, PA: Association for Computational Linguistics, 86–90.
- Balahur, Alexandra / M. Hermida, Jesús / Montoyo, Andrés (2011): Detecting implicit expressions of sentiment in text based on commonsense knowledge. Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA '11). Stroudsburg, PA: Association for Computational Linguistics, 53–60.
- Beavers, John (2011): On affectedness. In: *Natural Language & Linguistic Theory* 29.2, 335–370.
- Burchardt, Aljoscha (2008): Modeling Textual Entailment with Role-Semantic Information. Ph.D. Thesis. Saarland University, 2008.
- Cambria, Erik / Havasi, Catherine / Hussain, Amir (2012): SenticNet 2: A Semantic and Affective Resource for Opinion Mining and Sentiment Analysis. Proceedings of the Twenty-Fifth International Florida Artificial Intelligence Research Society Conference, 202–207.
- Deng, Lingjia / Choi, Yoonjung / Wiebe, Janyce (2013): Benefactive/Malefactive Event and Writer Attitude Annotation. In: Proceedings of the 51st Annual Meeting of the ACL (vol. 2). Stroudsburg, PA: Association for Computational Linguistics, 120–125.

- Deng, Lingjia / Wiebe, Janyce (2014): Sentiment propagation via implicature constraints. In: Proceedings of the 14th Conference of the European Chapter of the ACL. Stroudsburg, PA: Association for Computational Linguistics, 377–385.
- Dolbey, Andrew / Ellsworth, Michael / Scheffczyk, Jan (2006): Bioframenet: A domain-specific framenet extension with links to biomedical ontologies. KR-MED Conference 2006, 87–94.
- Ekman, Paul (1992): An argument for basic emotions. In: *Cognition & Emotion* 6 (3-4), 169–200.
- Elliott, Clark (1982): The affective reasoner: A process model of emotions in a multi-agent system. Ph.D. Thesis. Northwestern University, Evanston, Illinois.
- Fillmore, Charles J. (1976): Frame semantics and the nature of language. In: *Annals of the New York Academy of Sciences*, Volume 280, 20–32.
- Fillmore, Charles J. (1982): Frame semantics. In: *The Linguistic Society of Korea* (ed.), *Linguistics in the morning calm*. Seoul: Hanshin Publishing Corp., 111–137.
- Fillmore, Charles J. (1985): Frames and the semantics of understanding. In: *Quaderni di semantica* 6.2, 222–254.
- Hobbs, Jerry R. / Gordon, Andrew (2008): The Deep Lexical Semantics of Emotions. Proceedings of the International Conference on Language Resources and Evaluation (LREC) 2008. Workshop on Sentiment Analysis: Emotion, Metaphor, Ontology, and Terminology, Marrakech, Morocco, May 2008.
- Johnson-Laird, Philip Nicholas / Oatley, Keith (1989): The language of emotions: An analysis of a semantic field. In: *Cognition and emotion* 3 (2), 81–123.
- Lakoff, George / Johnson, Mark (1980): *Metaphors we live by*. Chicago: University of Chicago Press.
- Maks, Isa / Vossen, Piek (2011): A verb lexicon model for deep sentiment analysis and opinion mining applications. In: Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis. Stroudsburg, PA: Association for Computational Linguistics, 10–18.
- Martin, James R. (2003): Introduction. In: Macken-Horarik, Mary / Martin, James R. (eds.): *Negotiating heteroglossia: Social perspectives on evaluation*. Special issue of *Text – Interdisciplinary Journal for the Study of Discourse*, 23(2), 171–181.

- Martin, James R. / White, Peter R. R. (2005): *The language of evaluation*. Basingstoke / New York: Palgrave Macmillan.
- Ortony, Andrew / Clore, Gerald L. / Foss, Mark A. (1987): The referential structure of the affective lexicon. In: *Cognitive Science* 11 (3), 341–364.
- Ortony, Andrew / Clore, Gerald L. / Collins, Allan (1988): *The cognitive structure of emotions*. Cambridge: Cambridge University Press.
- Plutchik, Robert (1980): *Emotion: A psychoevolutionary synthesis*. New York: Harper & Row.
- Reschke, Kevin / Anand, Pranav (2011): Extracting contextual evaluativity. In: Bos, Johann / Pulman, Stephen (eds.): *Proceedings of the Ninth International Conference on Computational Semantics*. Stroudsburg, PA: Association for Computational Linguistics, 370–374.
- Ruppenhofer, Josef / Ellsworth, Michael / Petruck, Miriam R. L. / Johnson, Christopher R. (2010): *FrameNet II: Extended Theory and Practice*. Available at: <http://framenet2.icsi.berkeley.edu/docs/r1.5/book.pdf>.
- Ruppenhofer, Josef / Rehbein, Ines (2012): Semantic frames as an anchor representation for sentiment analysis. In: *Proceedings of the 3rd Workshop in Computational Approaches to Subjectivity and Sentiment Analysis*. Stroudsburg, PA: Association for Computational Linguistics, 104–109.
- Ruppenhofer, Josef (2013): Extending FrameNet for Sentiment Analysis. In: *Veredas* 17 (1), 66–81.
- Russell, James (1980): A circumplex model of affect. In: *Journal of Personality and Social Psychology* 39, 1161–1178.
- Schank, Roger C. / Abelson, Robert P. (1977): *Scripts, Plans, Goals and Understanding*. London: Lawrence Erlbaum Associates.
- Scherer, Klaus R (2000): Psychological models of emotion. In: Borod, J. C. (ed.), *The neuropsychology of emotion*, 137–162.
- Schmidt, Thomas (2009): The kicktionary – A multilingual lexical resource of football language. In: Boas, Hans C. (ed.): *Multilingual Framenets in Computational Lexicography*. Berlin / New York: de Gruyter, 101–132.
- Wiebe, Janyce / Deng, Lingjia (2014): *An Account of Opinion Implicatures*. arXiv preprint arXiv:1404.6491.
- Wierzbicka, Anna (1992): Defining emotion concepts. In: *Cognitive Science* 16.4, 539–581.