

# Characteristics of high agreement affect annotation in text

Cecilia Ovesdotter Alm  
Cornell University, USA  
cissioalm@gmail.com

## Abstract

The purpose of this paper is to present an unusual English dataset for affect exploration in text. It describes a corpus of fairy tales from three sources that have been annotated for affect at the sentence level. Special attention is given to data marked by high annotator agreement. A qualitative analysis of characteristics of high agreement sentences from H. C. Andersen reveals several interesting trends, illustrated by examples.

## 1 Introduction

Meaning is essential to language. The importance of expressive, attitudinal/emotive, or social/interpersonal meaning has been noted by prominent linguists (Bühler, 1934; Lyons, 1977; Jakobson, 1996; Halliday, 1996). However, affect is still an understudied phenomenon in linguistics, although many affective computing applications actually apply to language (Picard, 1997).

The motivation behind this discussion is to bring a special and rather unique dataset to the attention of researchers in the field of natural language processing, affective computing, and related areas. This paper discusses affect representation, presents an affect dataset, and then focuses on clear-cut cases of affective meaning and expression in text with a summary of an analysis of data for which human annotators highly agreed on the assignment of affect labels. For dataset results in supervised classification (including experimentation on high agreement data), cf. Alm (2009).<sup>1</sup>

## 2 Affect representation

Affect can be modeled, e.g. as categories (Ekman, 1994), dimensions (Osgood, 1969), by fo-

<sup>1</sup>For details on this dataset and experimentation conducted with it, readers should consult my book (Alm, 2009), which exceeds this paper in scope and depth.

cus on appraisal (Ortony et al, 1988), or on experience of physical and bodily responses (Cornelius, 2000). There is a lack of consensus on a model of affect (Picard, 1997; Scherer, 2003) and controversy surrounds such modeling. Pragmatically, different views of affect complement each other and jointly create a basis for understanding affective language phenomena. Affect modeling decisions are arguably application dependent. For a detailed literature review on previous work on how to characterize affect, affect in text-based linguistics and in subjective NLP or speech technology, and tales and oral narratives, see Alm (2009). Also see <http://emotion-research.net/>.

Resulting originally from an interest in text analysis for child-directed expressive text-to-speech synthesis, this dataset relies on a categorical annotation scheme of basic emotions; a model supported by the compelling observation that emotive facial expressions were cross-culturally recognized well above chance (Ekman and Friesen, 1998). In vision and speech research “the Big Six” (Cornelius, 2000) (i.e. *happiness, fear, anger, surprise, disgust, and sadness*) appear quite often. Nevertheless, the Ekmanian view remains controversial. For instance, Russel and Fernández-Dols (1998) have critiqued the relevance, methods, and rigor of the “Facial Expression Program” for emotion. One alternative is free labeling (i.e. annotators may come up with their own labels), but that may result in impractical, large label sets. A study grouping items from open-ended responses to a perception test on characterizing certain fairy tale sentences noted that although other cases occurred, Big Six emotions were frequent in answers (Bralè et al, 2005).

As regards the dataset’s use of affect categories, several empirical studies have shown above chance performance for recognition of categorical emotions in classification tasks involving prosody. Categorical labels may be more straightforward

for annotators to conceptualize compared to dimensional scales, as participants pointed out in a study (Francisco and Gervas, 2006). Also, categories are arguably suitable for pedagogy, and they naturally fit computational classification. A basic affect category is also broad enough to span related affect states, e.g. the *emotion family* (Ekman, 1994) of *angry* could also cover concepts such as *irritated*, *annoyed* and *enraged*.<sup>2</sup> Finally, the foundational nature of basic, categorical affects intuitively seems to fit a child-directed context and fairy tales contents, which may include certain canonical topics and behaviors, compared to more spontaneous discourse.<sup>3</sup>

### 3 Corpus data overview

The affect dataset consists of 176 stories (more than 15,000 sentences) by Beatrix Potter, the Brothers Grimm and H. C. Andersen, manually annotated at the sentence level by pairs of annotators.<sup>4</sup> For the annotation process, annotators read tales and had to make a choice from a set of affect categories for sentences. Each sentence was given four affect labels since each of two annotators assigned both a *primary emotion* (guided by the presence of a *feeler*, mostly a character or character type in the text) and a background *mood* to a sentence. The four labels were then combined into a sentence’s affect labels. For more details on the annotation process, cf. (Alm, 2009). The label set consisted of a set of categorical affect labels. Prior to the analysis below, ANGRY and DISGUSTED were merged (motivated by data sparsity and related semantics) into one category, as were POSITIVELY and NEGATIVELY SURPRISED, yielding a merged set of affect labels: ANGRY-DISGUSTED, FEARFUL, HAPPY, NEUTRAL, SAD, SURPRISED.

Interannotator agreement can be an artifact of annotation scheme and procedure. For example, pairs might be trained to annotate similarly, across-the-board rules (e.g. *questions* are negative) might ignore subtle decisions, or problematic items might be removed. Such approaches may yield higher agreement, cleaner data, and perhaps better performance and more consistent

<sup>2</sup>Categories do not exclude adding intensity for approximating an arousal dimension, arguably relevant for speech.

<sup>3</sup>Naturally, tales also encompass narrative complexity.

<sup>4</sup>The annotated data are available at the author’s website (both the full dataset and the high agreement subsets). For instance, for the high agree affect data, a storyname is followed by its corresponding high agree affective sentences in the following format: sentence-id-in-story@label-code@sentence.




Figure 1: (Dis)agreement: merged labels

trained applications. But, the relevance of that for study of linguistic behavior is less clear. Zaenen (2006) noted that “[f]or interannotator agreement, it suffices that all annotators do the same thing. But even with full annotator agreement it is not sure that the task captures what was originally intended” (577); this should not be confused with understanding a linguistic issue. Fig. 1 reports on a diagnostic alternative with the ratios of (dis)agreement types. This avoids the concept of *ground truth*, which may not hold for all language phenomena. Affect, which is highly subjective, is arguably better captured by flexible *acceptability*.<sup>5</sup>

Fig. 1 shows that sentences only labeled NEUTRAL were frequent, as were disagreements, which were more common for sentences marked both with NEUTRAL and one or more affect classes. This parallels findings for polarity expressions in subjective texts (Wilson et al, 2005), and shows that the border between affective and neutral is fuzzy. (Affect perception lacks clear definitions and is subjective, and neutrality suffers from the same dilemma.) A sentence with *high agreement* affect was defined as all four primary emotion and mood labels having the same *affective* label (given the merged label set). These were more common than mixed affective labels.

### 4 High agreement in H. C. Andersen

This section examines the subset of high agreement sentences in the H. C. Andersen data from a qualitative-interpretive perspective. The analysis is not intended as rigid categorization, but rather to get an overall idea of why high agreement might occur on affect labels across annotators. Isolated sentences were extracted and mostly examined that way, rarely considering context. This

<sup>5</sup>Regular agreement scores for the corpus would be low.

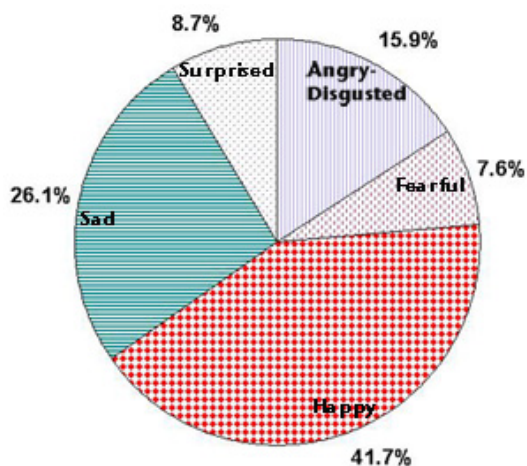


Figure 2: Distribution of 460 H. C. Andersen high agreement affective sentences across affect labels focused the analytical scope.<sup>6</sup> Five annotators engaged with the overall H. C. Andersen subcorpus of 77 tales. 460 sentences were marked by affective high agreement, given the five affective classes. The distribution of affective classes for this subset is in Fig. 2, with HAPPY and SAD being most frequent.

#### 4.1 Characteristics: high agreement affect

The below overview lists characteristics observed in an analysis on the H. C. Andersen high agreement data. It briefly describes each characteristic and lets an example illustrate it. For more discussion, examples, word lists etc., see Alm (2009). The characteristics occur in some and not all sentences; some frequently, others more rarely. Often, several jointly characterize a sentence.

The illustrative sentence examples in this section use the following format: Affect labels are in small caps and sentences are in italics. Also, phrases in bold-face illustrate the discussed characteristic, whereas phrases that annotators noted are underlined (single underscore for non-overlapping vs. double underscore for overlapping mark-up), and their *feeler/s* for the primary emotion annotation is/are included (with annotator subscripts to show if they had indicated the same or not) in parenthesis in small caps.

##### 4.1.1 Affect words

Content words that directly name an affective state (e.g. reflecting a particular intensity) are common in high agreement sentences, cf.:

<sup>6</sup>Annotators' noted *feeler* and emotional/connotative phrases for the sentences were inspected.

ANGRY-DISGUSTED: *They buzzed round the prince and stung his face and hands; **angrily** he drew his sword and brandished it, but he only touched the air and did not hit the gnats.* (VILLAIN<sub>1,2</sub>)

That narration can directly announce affective states is an indication of the important narrative role affect can play in stories. Also, Wilson and Wiebe (2003) interestingly noted that annotators agreed more strongly with strong subjective expressions, which affect words are examples of. Some illustrative affect words from the examined data are (for SURPRISED): *alarmed, astonished, astonishment, shocked, shocking, startled, surprised*. Special cases include *negation* (e.g. *not happy* for SAD); *figurative/idiomatic phrases* (e.g. *one of his heartstrings had broken* for SAD); or appearance with more than one affect (e.g. *anguish* for SAD or FEARFUL).

##### 4.1.2 Words for related/contrastive affect states

Expressions in the sentential context naming related or contrastive affective states not in the label set (e.g. *dull, pride, relief, or shame*) may also help evoke a particular affect, as in:

HAPPY: *They looked at Little Claus ploughing with his five horses, and he was so **proud** that he smacked his whip, and said, "Gee-up, my five horses."* (HERO<sub>1,2</sub>)

##### 4.1.3 Affect related words or expressions

Lexical items or phrases which describe actions, properties, behaviors, cognitive states, or objects associated with particular affects occur frequently in the examined high agreement subset, e.g. as in: HAPPY: *They **laughed** and they **wept**; and Peter **embraced** the old Fire-drum.* (HERO<sub>1</sub>, (TRUE) MOTHER<sub>2</sub>, (TRUE) FATHER<sub>2</sub>)

Some more prominent affect related lexical items include *weep, kiss, laugh, cry (= weep)*, and forms of *pleasure, tears, and smile*. Expressions of weeping or tears often appear with sadness, but may also depict happiness. *Negations* may occur.

##### 4.1.4 Polarity words and expressions

Words or expressions of positive or negative polarity can help to set the scene with a particular affective mode, in particular with relation to context and acquired knowledge. Expressions of opposing polarity may be used as a contrast, as in:

HAPPY: *It became a **splendid** flower-garden*

to the *sick* boy, and his *little treasure* upon earth. (SICK BOY<sub>1,2</sub>)

Modifiers can intensify the affective load. Lexical words and phrases may have *permanent* vs. *occasional* attitudinal meaning (Hedquist, 1978).

#### 4.1.5 Knowledge and human experience

Readers may from experience associate acquired knowledge about situations, visualizations, and behaviors with particular affects. For example, it is common knowledge that starving is traumatic: SAD: *He was **hungry and thirsty**, yet no one gave him anything; and when it became dark, and they were about to close the gardens, the porter **turned him out**.* (HERO<sub>1,2</sub>).

Story worlds tend to involve canonical representations of characters, actions, functions, situations and objects. Surrounding *context* can be important for affective interpretations. Scenarios may include, e.g. *an inspiration from weather, flowers, nature, or God; singing (or dancing, jumping); physical lack and need; sleep deprivation or allowance; addiction; incapability; unexpected observation; appearance/posture (or intonation); contextual guidance; or relate to marriage* (see (Alm, 2009) for examples). In fact, arguably most discussed characteristics can be traced to acquired knowledge, experience, associations, or context.

#### 4.1.6 Speech acts

Speech acts reflect a certain kind of communicative knowledge that can have affective meaning (such as *cursing, insulting, commanding*), e.g.:

ANGRY-DISGUSTED:

*Let her be expelled from the congregation and the Church.* (VILLAIN<sub>1,2</sub>)

#### 4.1.7 Types of direct speech

Direct speech may be used by characters in tales to express affect. This might include *speaking excitedly, (WH)-exclamations* or *(WH)-questions, short utterances, interjections* (and *sound effects*), such as *ah, alas, hurrah, o God, sorry, thump, ugh*. Direct speech can be introduced by words of speaking, as in:

FEARFUL: *“Mercy!” cried Karen.* (HEROINE<sub>1,2</sub>)

#### 4.1.8 Mixed emotions

Affective high agreement sentences also include cases of mixed emotions, e.g. affect or affect-related words referring to more than one affect. The ‘winning’ affect may be inferred. Contrast

might make it more prominent, as in:

HAPPY (mixed SAD): *He now **felt glad** at having **suffered sorrow and trouble**, because it enabled him to **enjoy so much better all the pleasure and happiness** around him; for the great swans swam round the new-comer, and stroked his neck with their beaks, as a welcome.* (MAIN CHARACTER/HERO<sub>1,2</sub>)

#### 4.2 Tendencies of particular affect categories

Lastly, there may be trends for particular characteristics associating more or less with a particular affect. For example, in this subset, FEARFUL sentences seem often to contain affect or affect related words, whereas SURPRISED sentences may quite often be characterized by various types of direct speech or involve unexpected observations.

### 5 Conclusion

This paper brought attention to an affect dataset, and discussed (mostly surface) characteristics in its H. C. Andersen high agreement subset, illustrating the complexity of affect cues, without claiming an exhaustive analysis. It also tentatively hypothesized that some characteristics may show particular affinity with certain affects.

The high agreement sentence data may be particularly interesting for affect research, while other parts of the annotated, larger corpus may reveal insights on affect variation in text and perception thereof (bearing in mind that the dataset is not necessarily representative across domains and text types, nor of contemporary texts).

Lastly, as noted above, developed knowledge, experience, associations, and context appear very important for affect understanding. This is also a substantial part of what makes the problem of automatically predicting affect from text so challenging; it involves levels of deep cognitive understanding rather than just extractable surface features. Whereas the discussed characteristics naturally do not constitute the answer to affect understanding, they may inform future search for it. Deep understanding and continuous, as opposed to static, computational development of affective understanding remain crucial areas of future work for expressive NLP applications.

### Acknowledgments

Thanks to R. Sproat, R. Proaño, and reviewers. Project funded by NSF (award ITR-#0205731).

## References

- Alm, Cecilia Ovesdotter. 2009. *Affect in Text and Speech*. VDM Verlag: Saarbrücken.
- Bralè, Véronique, Valérie Maffiolo, Ioannis Kanellos, and Thierry Moudenc. 2005. Towards an expressive typology in storytelling: A perceptive approach. In Jianhua Tao, Tieniu Tan, and Rosalind W. Picard (Eds.), *Affective Computing and Intelligent Interaction, First International Conference, ACII 2005, Beijing, China, October 22-24, 2005, Proceedings*, 858-865.
- Bühler, Karl. 1934. *Sprachtheorie: Die Darstellungsfunktion der Sprache*. Stuttgart: Gustav Fischer Verlag.
- Cahn, Janet E. 1990. The generation of affect in synthesized speech. *Journal of the American Voice I/O Society* 8, 1-19.
- Cornelius, Randolph R. 2000. Theoretical approaches to emotion. In *Proceedings of the ISCA Workshop on Speech and Emotion*, 3-10.
- Ekman, Paul. 1994. All emotions are basic. In P. Ekman and R. J. Davidson (Eds.), *The Nature of Emotion: Fundamental Questions*. Oxford: Oxford University Press, 15-19.
- Ekman, Paul and Wallace V. Friesen. 1998 [1971] Constants across culture in the face and emotion. Jenkins, Jennifer M and Oatley, Keith and Stein, Nancy L. (eds). *Human Emotions: A Reader*. Malden, Massachusetts: Blackwell, 63-72.
- Francisco, Virginia and Pablo Gervás 2006. Exploring the compositionality of emotions in text: Word emotions, sentence emotions and automated tagging. In *AAAI-06 Workshop on Computational Aesthetics: Artificial Intelligence Approaches to Beauty and Happiness*.
- Halliday, Michael A. K. 1996. Linguistic function and literary style: An inquiry into the language of William Golding's *The Inheritors*. Weber, Jean Jacques (ed). *The Stylistics Reader: From Roman Jakobson to the Present*. London: Arnold, 56-86.
- Hedquist, Rolf. 1978. *Emotivt spåk: En studie i dagstidningarnas ledare*. Ph.D. Thesis. Umeå.
- Jakobson, Roman. 1996. Closing statement: Linguistics and poetics. Weber, Jean Jacques (ed). *The Stylistics Reader: From Roman Jakobson to the Present*. London: Arnold, 10-35.
- Lyons, John. 1977. *Semantics* volumes 1, 2. Cambridge: Cambridge University Press.
- Ortony, Andrew, Gerlad L. Clore, and Allan Collins. 1988. *The Cognitive Structure of Emotions*. Cambridge: Cambridge University Press.
- Osgood, Charles E. 1969. On the whys and wherefores of E, P, and A. *Journal of Personality and Social Psychology* 12 (3), 194-199.
- Picard, Rosalind W. 1997. *Affective computing*. Cambridge, Massachusetts: MIT Press.
- Russell, James A. and José M. Fernández-Dols 1998 [1997]. What does a facial expression mean? Jenkins, Jennifer M and Oatley, Keith and Stein, Nancy L. (eds). *Human Emotions: A Reader*. Malden, Massachusetts: Blackwell, 63-72.
- Scherer, Klaus R. 2003. Vocal communication of emotion: A review of research paradigms. *Speech Communication* 40 (1-2), 227-256.
- Wilson, Theresa, Janyce Wiebe, and Paul Hoffman. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. *Proceedings of HLT/EMNLP*, 347-354.
- Wilson, Theresa and Janyce Wiebe. 2003. Annotating opinions in the world press. *4th SigDial workshop on Discourse and Dialogue*.
- Zaenen, Annie. 2006. Mark-up barking up the wrong tree. *Journal of Computational Linguistics* 32 (4), 577-580.