

Do Happy Words Sound Happy?

A study of the relation between form and meaning for English words expressing emotions

Vivi Nastase
EML Research gGmbH
Heidelberg, Germany
nastase@eml-research.de

Marina Sokolova
Université de Montreal
Montreal, Quebec, Canada
sokolovm@iro.umontreal.ca

Jelber Sayyad Shirabad
University of Ottawa
Ottawa, Ontario, Canada
jsayyad@site.uottawa.ca

Abstract

This paper presents a study of the relation between a word’s form and the emotion it expresses. We analyze the possibility that the form of words expressing emotions is not completely arbitrary, but in fact, their sound evokes the emotion conveyed. We explore the relation between word form and emotions using a variety of word form representations and machine learning methods. We first show that words expressing an emotion are more similar among them than with words expressing other emotions, and then we discuss the sounds of emotions.

Keywords

sentiment analysis, word form

1 Introduction

A word has two components: **word form**, a sequence of sounds (pronunciation) and, possibly, letters/characters (written form), and **meaning**. The word form is also called **signifier**, and its meaning, or referent in the world, is called **signified**: the word form *tree* with the pronunciation $/trE/$ ¹ has as referent in the real world a TREE entity.² While it is usually accepted that the relation between signifier and signified is largely arbitrary [5], the idea that sounds may carry meaning has appeared at several points in time [8], and is still a matter of debate and research.

In this paper we study the relationship between signifier and signified for a class of words which can be particularly susceptible to the way a word sounds: words that express emotions – either positive or negative, or a more fine grained range (anger, disgust, fear, joy, sadness, surprise). If there is such a relationship, we should be able to build a model of the word forms expressing different emotions, and use this model to predict with good accuracy the emotions expressed by unseen words.

Apart from a purely theoretical benefit, finding a relation between the way the words sound and the emotion expressed contributes to research in sentiment analysis, very much part of the highly explored areas

of NLP these days, authorship analysis and other research areas.

We work with data annotated with emotion tags: WordNet Affect and the dictionary from the Linguistic Inquiry and Word Count system. We work with the pronunciation and written form of a word. We represent the word form in various ways, using separately the written and pronunciation versions. Results of our experiments show that indeed, there is a connection between word forms and the emotion conveyed. Happy words sound more like other happy words than other emotion words. Whether the sounds of happy words are indeed happy sounding is a harder question to answer, as perception is subjective. We discuss the sounds of emotions based on the most salient features in our experiments and analysis in the literature.

There is certainly more to the word form than just the sequence of sounds or letters (e.g., variations in tone, stress, cadence) – imagine someone saying in a sarcastic tone “I’m fine”, as opposed to someone saying in a merry voice the same words. Any information coming from speech analysis will be a bonus, ready to be added on top of our findings from experiments with the form only.

2 Motivation

It is a long held belief that the association between a word-form and its meaning is arbitrary [5]: there is nothing about a TREE that evokes the sequence of letters or sounds that form the English word *tree*. Support of this theory comes from language variation: a TREE is called *tree* in English, but *Baum* in German, *albero* in Italian, and numerous other variants in the languages of the world. If there was anything intrinsic to TREE that would link it to the form *tree*, it would have been called the same in all languages.

There are also onomatopoeic words, which sound like the concept they are describing [2]. Onomatopoeia are language specific. In English lions *roar*, cats *purr*, flies *buzz*, snakes *hiss*, fireworks go *boom* and *bang*.

In between the two extremes of total arbitrariness of form relative to meaning and identity of the two, there are *mellifluous* words. Coming from the Latin *mellifluus* = *mel*(honey)+*fluere*(flow) – dripping with honey – mellifluous has come to refer to words whose sounds evoke the concepts they refer to. Such words were particularly exploited for effects in poetry [16]. We also use them in our everyday speech: we *hush*

¹ From the online version of the Merriam-Webster: <http://www.m-w.com>.

² For the remainder of the paper, the signifier will be written in *italics*, and the signified in SMALLCAPS.

to make silence, we *mumble* when we speak in a low inarticulate manner.

Arbitrariness of the connection between sound and meaning is not universally accepted. The theory of *sound symbolism* or *phonosemantics*, according to which most words in a language fall into a category similar to *mellifluous* – every sound carries a certain meaning, which evoke certain aspects of a concept whose name contains this sound – has ancient roots. Plato, through his characters in the *Cratylus* dialogue – Hermogenes and Socrates – discusses the provenance of words. Socrates proposes that there is a connection between the way words sound and their signifiers. As an example, he gives the greek letter ρ (rho), which for him expresses motion. A number of (Greek) words containing ρ are brought up in support of this hypothesis, for which Hermogenes provides afterwards a plethora of counter-examples.

The idea that sounds carry meaning has reappeared throughout history. Locke’s *An Essay on Human Understanding* (1690) counters this idea. Leibniz’s book *New Essays on Human Understanding* (1765) critiques Locke’s essay. Leibniz proposes a moderate view, in which words and their referents are neither related by perfect correspondence, nor by complete arbitrariness. A detailed history of phonosemantics is presented by Genette [8], and a historical review plus recent research and developments are presented by Magnus [15].

An interesting view on the relation between sound and meaning, and the possible connection between the two, is proposed by Jakobson [11]. In Lecture VI he says: “The intimacy of connection between the sounds and the meaning of a word gives rise to the desire of speakers to add an internal relation to the external relation, resemblance to contiguity, to complement the signified by a rudimentary image”. In other words, Jakobson says that resemblance between sound and meaning is in the ear and mind of the beholder. This may lead to a “natural selection” of words, based on the way they sound, as suggested by Otto Jespersen: “There is no denying that there are words which we feel instinctively to be adequate to express the ideas they stand for. ... Sound symbolism, we may say, make some words more fit to survive.” [12]. Firth [7] and Sapir [19] also share such a middle-ground view of sound symbolism. In their view, speech sounds carry meaning, but rather than being inherent to them, it is a result of what Firth called “phonetic habit”, “an attunement of the nervous system”.

3 Signifier and signified

We set out to investigate the connection between the signifier, or word form, and signified, or meaning, for the English words that express emotions. Because we propose that words expressing emotions are mellifluous words, we do not seek a relation between form and exact meaning, but rather form and some aspect of the meaning - in our case, the emotion conveyed.

The signifier The signifier, in our case, can have a written and a spoken form. A TREE is called */trE/* and written *tree* in English. The pronunciation is a sequence of sounds (phonemes). According to research in speech analysis, phonemes are not the smallest units of speech. Individual phonemes can be represented

Phoneme	Example	Transcription
AA	alarm	AH0 L AA1 R M
AE	amorous	AE1 M ER0 AH0 S
CH	charm	CH AA1 R M
EH	enchant	EH0 N CH AE1 N T
T	tickle	T IH1 K AH0 L
Y	euphoria	Y UW0 F AO1 R IY0 AH0

Table 1: A sample of phonemes, words and their phonetic transcription

through values of a set of parameters, or features, that capture the configuration of the vocal tract that produces each sound and other acoustic features. We investigate each of these three variants of representing a word form.

letters : In English words are not pronounced as they are written. However, the way words are spelled may be closer to the words’ etymological roots than their pronunciation is. As an example, the word *delight*, comes from the Old French word *delit*, *delitier* which in turn comes from the Latin *delectare*³. The letter *e* in *delight* is pronounced /i/ as in *bit*, while in its etymological roots, it is pronounced /e/ as in *bet*. Since texts are more readily available than word pronunciations, this type of word form is also the easiest to analyze.

pronunciation : Pronunciation of letters in English, especially vowels, depends on their context. Dictionaries provide a transcription of words into their phonetic equivalent. In this representation, each sound (which may correspond to one or more of a word’s letters) is represented by a special symbol. We use CMU’s pronunciation dictionary⁴, which contains approximately 125,000 words and their transcriptions. The transcriptions’ “alphabet” consists of 39 phonemes, and three extra digits for primary and secondary stress information (0 - no stress, 1 - primary stress, 2 - secondary stress). A sample of phonemes and word pronunciations are presented in Table 1.

phonetic-features : The sounds of language can also be described in terms of *phonological features* – “configurations” of the vocal tract and acoustic characteristics. From the existing phonological feature systems – [13], [9], [3] – we use the Sound Pattern of English (SPE) [3].

SPE consists of 14 binary features, which describe the tongue body position (high, back, low), tongue tip position (anterior, coronal), lips’ configuration (round), configurations affecting the air flow – by constriction, vibration of vocal folds or blocking with the tongue or lips (tensed, voiced, continuant, nasal, strident) and acoustic characteristics (vocalic, consonant, silence). Examples of phones with their SPE representation are shown in Table 2.

The signified The signified component of our data comes from emotion tags, from the two sets: { *anger*,

³ From the Online Etymology Dictionary: <http://www.etymonline.com>.

⁴ The CMU pronunciation dictionary is freely available at <http://www.speech.cs.cmu.edu/cgi-bin/cmudict> . We have used version 0.6d.

	v	c	h	b	l	a	c	r	t	v	c	n	s	s
	o	o	i	a	o	n	o	o	e	o	o	a	t	i
	c.	n	g	c	w	t.	r.	u	n	i	n	s	r	l
	s.	h	k				n.	s.	c.	t.	a.	i.	e.	
ae (bat)	+	-	-	+	+	-	-	-	+	+	+	-	-	-
b (bee)	-	+	-	-	-	+	-	-	-	+	-	-	-	-
iy (beet)	-	-	+	-	-	-	-	-	+	+	+	-	-	-
m (mom)	-	+	-	-	-	+	-	-	-	+	-	+	-	-
ow (boat)	+	-	-	+	-	-	-	+	+	+	+	-	-	-
sh (she)	-	+	+	-	-	-	+	-	-	-	+	-	+	-

Table 2: Examples of sound representation using the SPE system

disgust, fear, joy, sadness, surprise } , or { *positive, negative* } . Our hypothesis is that words expressing emotions are mellifluous words – that is, the way the sound evoke the emotion they convey. Psychological research proposes the following basic emotions: { *anger, disgust, fear, joy, sadness, surprise* } [6]. We study whether analysis of word form allows us to predict whether the word expresses one of these basic emotions. Because much research in the domain of sentiment analysis works at a coarser level of emotions – *positive* and *negative* – we also study the relation between word forms and these broader emotion categories.

4 Emotion-tagged words

Assigning an emotion tag to words is not an easy task. Potentially, for any word one may perceive an emotional dimension, either directly from the word’s meaning, or through the word’s associations with emotionally charged words or situations.

The words we are most interested in are words that actually express an emotion, such as *happy, joy*. We focus on WordNet Affect [21] and LIWC [18] data because they contain words that express emotions, rather than having a semantic orientation. The word *knowledge* for example, does not have an emotion tag in WordNet Affect, but it has a positive tag in the General Inquirer data. Other resources include the General Inquirer data⁵ and the list of positive and negative adjectives used by Hatzivassiloglou and McKeown [10]⁶.

WordNet Affect WordNet-Affect is an extension of WordNet with affective tags. Words that have an EMOTION tag, were recently more fine-grained reannotated: { *joy, fear, anger, sadness, disgust, surprise* } [21]. The choice for the six emotions comes from psychological research into human (non-verbally expressed) emotions [6].

Linguistic Inquiry and Word Count Linguistic Inquiry and Word Count (LIWC) focuses on the analysis of text and computing statistics along 82 dimensions, such as positive emotions, negative emotions, present, future, space, motion, occupation, physical, metaphysical, body [18]. We use the words

⁵ The General Inquirer lexicon is freely available for research purposes from <http://www.wjh.harvard.edu/inquirer/>.

⁶ These and other sentiment annotated resources are available from Janyce Wiebe’s web site <http://www.cs.pitt.edu/wiebe>

Resource	class	count	with pron.
WordNet Affect	anger	240	101 (25%)
	disgust	48	17 (4.2%)
	fear	134	49 (12.13%)
	joy	364	152 (37.63%)
	sadness	187	59 (14.6%)
	surprise	70	26 (6.44%)
	total	1043	404 (100%)
LIWC	negative	345	283 (59.21%)
	positive	265	195 (40.79%)
	total	610	478 (100%)

Table 3: Word-sentiment counts from WordNet Affect and LIWC

with positive and negative emotion dimensions.

Table 3 contains information about the number of unique words for each of the WordNet affect and LIWC emotions. Column 3 shows the word count for each emotion, and column 4 shows the word count after filtering morphologically related words and after looking up the word’s pronunciation in the CMU dictionary. The experiments are run only using words that have a pronunciation in the dictionary, to allow for comparison of performance for the different representations. We filtered morphologically related words by performing stemming (using Porter’s Stemmer), an extra step of cutting off suffixes (such as *-fully, -ful, -some, -ness*) to catch words with multiple suffixes, and finally word matching. We also eliminate words with the suffix “less”, because the emotion the word stem expresses and the emotion expressed by the full word are different. Words with negative prefixes (un-, in-) are kept, because it is harder to detect whether a starting sequence *un* or *in* is actually a prefix or not. Also, the bigram representation, discussed below, will cover these prefixes (as opposed to the suffix *-less* for which a 4-gram representation would be necessary).

5 Learning experiments

Our working hypothesis is that word forms expressing the same emotion share sound/pronunciation characteristics – in other words, they sound similar in certain ways. The similarities may be at the smallest level – letter, sound, sound feature – or at a more complex level – letter or sound sequences, combinations of sound features. We build data representations at these three levels, and test out hypothesis using decision tree (J48, ADTree⁷) and memory based (IBK) algorithms in Weka [22], in 10-fold cross-validation experiments.

Data representation Following these considerations, we have produced a series of representations for the data, which vary along two dimensions: analyzed unit (unigrams and bigrams) and unit representation (letter, pronunciation and sound features).

We split each word into three segments – beginning unit, ending unit, and the middle segment which contains everything in-between. For example, if we consider the word *admire* with a bigram letter representation, it will have the following segments: beginning – *ad*, middle – *dmir*, end – *re*. Each word is represented

⁷ We use Weka’s MultiClassClassifier to perform multi-class classification with ADTree.

Repres.	# of Features	Repres.	# of Features
W1Let	71	W2Let	498
W1P	126	W2P	711
W1C	42	W2C	588
L1Let	68	L2Let	498
L1P	123	L2P	731
L1C	42	W2C	588

Table 4: Number of features used in each representation method

in terms of features for each of these three segments. For each segment, the features represent aggregated statistics for the units in this segment. For letter feature a in the middle segment, for example, the value is the number of occurrences of a in the middle segment.

For the bigram letter representation of the word *admire*, the following features will have non-zero values: for the beginning segment – *ad*, for the middle segment – *dm*, *mi*, *ir*, for the end segment – *re*.

Table 4 shows the number of features for each data set generated for the 6 possible variations. In the table, and in the discussion that follows, we will use the abbreviations: *data sets*: WordNet Affect (W), LIWC (L); *units*: unigram (1), bigram (2); *unit representation*: letters (Let), pronunciation/phonetic (P), SPE codes (C) levels.

For letter- and pronunciation/phonetic-based representation, the features are determined by the n -gram letter and phoneme sequences that actually appear in our list of words. For phonological features we consider the set of 14 SPE features for each word segment (beginning, middle, end). Each feature has numeric values, and counts the number of time it appears in the corresponding segment of the word (beginning, middle, end). All features are numeric and their value is the number of occurrence of the feature (e.g. letter a or phoneme IY) in the corresponding segment of the word. The phonetic-based representation contains two extra features – for primary and secondary stress, as indicated in the pronunciation dictionary. These two features take as value the phoneme that was stressed (always corresponding to a vowel).

Results for WNAffect data We evaluate the quality of classification by computing the average accuracy (Acc) and average precision (P), recall (R) and F1 score (F) for each emotion class in 10-fold cross validation experiments. Our experiments have shown that dropping the features for the end of word segment has a positive impact on performance.

The best results, in terms of accuracy for the WNAffect data were obtained with the unigram letter-based (39.85%) and unigram code-based (38.11%) unigram representation, on word representations based only on the beginning and middle segment. In this 6-class learning problem the baseline accuracy is 37.63%, corresponding to classifying everything as *joy*, the majority class.

Table 5 shows the best results in terms of F-score for each class (emotion) in the WordNet Affect data, also in the multi-class learning setting. For detailed results on each emotion class we use a baseline which guesses the class with a distribution that matches the one in the data set (this balances precision and recall). The baseline F-score values are given by the distribution presented in Table 3, repeated here on row 2.

Method	anger	disgust	fear	joy	sadness	surprise
baseline	25%	4.2%	12.13%	37.63%	14.6%	6.44%
IB1-1Let	40.9%	33.3%	33.3%	49.3%	33%	9.8%
IB1-1C	42.2%	16.2%	29.5%	45.1%	31.5%	26.7%
highest	44.8%	36.8%	34.5%	55%	33%	26.7%
values	IB2-1C	IB2-1Let	IB2-1C	IB2-1C	IB1-1Let	IB1-1C

Table 5: F-score results on 6-class classification into WordNet Affect emotions

The best recognized emotion from WordNet Affect’s emotion classes was *joy*. Despite variation in P , R , and F values for different representations and learning algorithms, *joy* was consistently the best classified emotion. Part of this may be due to the fact that it had the most examples (37.63%). The results show a statistically significant improvements over the baselines considered with a confidence of 0.05 with Weka’s two tailed test.

Results on LIWC data A selection of the most balanced results (in terms of F-score) for the LIWC data are presented in Table 5. The baseline F-score values are given by the distribution presented in Table 3. The performance increase over the baseline is statistically significant with a confidence of 0.05 (with Weka’s two tailed test). For this binary classification experiment, the baseline accuracy is 56.59%, corresponding to classifying everything as *negative*, the majority class. The best results, in terms of accuracy for the LIWC data are 62.3% (IBK, $k=55$, 1C) and 61.5% (J48, 2Let).

Method	positive	negative
baseline	40.79%	59.21%
IB1-2Let	45.5%	68.2%
IB1-1P	44.1%	67.9%
highest	45.5%	74.9%
values	IB1-2Let	IBK,K=55 -1C

Table 6: F-score results on binary classification on LIWC

For the LIWC data it is the negative emotion class that outperforms the positive one. The negative emotion also has more instances. Note that IBK ultimately outperforms other classifiers in finding balanced results for the binary problem.

6 The sounds of emotions

We have set out to study the connection between the way a word *sounds*, and the information conveyed. The learning experiments we performed show that words expressing the same emotion have more in common with each other than with words expressing other emotions. But do they sound as the emotion they convey? Does *happy* really sound happy?

Recognizing emotions in human speech is an intensely studied topic. The type of data used are recordings of utterances, whose sound signal is represented through a variety of features (such as pitch, energy, tone contour) [20],[4], [17]. Lee et al. [14] introduce five broad phoneme classes – vowel, stop, glide, nasal, fricative – to help in classifying utterances into 4 classes – angry, happy, neutral and other.

In learning experiments using Hidden Markov Models, they note that using phoneme classes in addition to the more traditional signal features leads to better emotion recognition. In particular, vowel sounds are good emotion indicators, and furthermore different vowels have different effects, possibly because of articulatory constraints: “less constricted low vowels such as /AA/ show greater effects than do high vowels like /IY/”. There are no details as to which vowels are predictive of which emotion class.

7 Conclusion

We have investigated the phonological properties of word-forms, to learn whether we can automatically predict the emotion a word expresses based on various representations of its form. The results show that all the representations used – word spelling, pronunciation, phonetic features – are useful for determining that words expressing the same emotion are alike in certain ways. The fact that different representations worked best for different emotion classes leads us to believe that with an ingenious combination of representations and feature selection we can improve on these results.

These results answer half of the question we had set out to investigate – whether words sound like the emotion conveyed. The other half is whether what happy words have in common is what makes them sound happy. To address this question we have found parallels between the features in our data representation that best discriminate among emotion classes, and speech characteristics identified through research in speech emotion recognition that correlate with the emotion conveyed.

We plan to experiment with alternative word representations – such as syllables, which are considered the phonological “atoms” of words – and to determine which part of the word is most expressive from the point of view of the emotion conveyed. Research based on the words’ etymological roots may show us if the link between form and meaning gets stronger as we go back in time. Next step is to expand the study to languages other than English, and to longer text units, such as blogs. In speech emotions are detectable, and the speaker conveys these through tone and other prosodic features. It would be interesting to see whether we can detect “sub-word” level features to detect emotions in blogs.

The degree of centrality of a word to the category of sentiment varies, and words expressing emotions can be approached from a fuzzy set point of view [1]. It would be interesting for us to study whether words more prototypical for the six basic emotions have stronger relations with their form than those towards the margins of the fuzzy set.

Compared to previous work in determining the semantic orientation of a word, our experiments rely on local information only – the word and its pronunciation – which simplifies much the data representation and has a light computation load. The type of representation we used may be added to collocation based features, to provide additional clues about the sentiment dimension of words.

Acknowledgments Partial support for this work is given by RALI (Guy Lapalme et al.), Université de Montreal. We thank the three anonymous reviewers for their great, insightful comments.

References

- [1] A. Andreevskaia and S. Bergler. Mining wordnet for a fuzzy sentiment: Sentiment tag extraction from wordnet glosses. In *Proceedings the 11th EACL-06*, pages 209–216, Trento, IT, 2006.
- [2] H. Bredin. Onomatopoeia as a figure and a linguistic principle. *New Literary History*, 27(3):555–569, 1996.
- [3] N. Chomsky and M. Halle. *The Sound Pattern of English*. MIT Press, 1968.
- [4] R. Cowie, E. Douglas-Cowie, N. Tsapatsoulis, G. Votsis, S. Kollias, W. Fellenz, and J. Taylor. Emotion recognition in human-computer interaction. *IEEE Signal Processing Magazine*, 18(1):32–80, 2001.
- [5] F. de Saussure. *Cours de linguistique générale*. Harrassowitz, Wiesbaden, 1916.
- [6] P. Ekman. An argument for basic emotions. *Cognition and Emotion*, 6:169–200, 1992.
- [7] J. R. Firth. Modes and meaning. In *Papers in linguistics 1934-1951*. Oxford University Press, London, 1951.
- [8] G. Genette. *Mimologiques: voyage en Cratylie*. Seuil, Paris, 1976.
- [9] J. Harris. *English Sound Structure*. Blackwell, 1994.
- [10] V. Hatzivassiloglou and K. McKeown. Predicting the semantic orientation of adjectives. In *Proceedings of the 35th ACL/8th EACL*, pages 174–181, Madrid, Spain, 1997.
- [11] R. Jakobson. *Lectures on Sound and Meaning*. MIT Press, Cambridge, MA, 1937.
- [12] O. Jespersen. *Language - its Nature, Development and Origin*. George Allen & Unwin Ltd., London, 1922.
- [13] J. King and P. Taylor. Detection of phonological features in continuous speech using neural networks. *Computer speech and language*, 14(4):333–353, 2000.
- [14] C. M. Lee, S. Yildirim, M. Bulut, A. Kazemzadeh, C. Busso, Z. Deng, S. Lee, and S. Narayanan. Emotion recognition based on phoneme classes. In *Proc. of ICSLP*, 2004.
- [15] M. Magnus. *What’s in a Word? Studies in Phonosemantics*. PhD thesis, University of Trondheim, Trondheim, Norway, 2001.
- [16] M. Nänny and O. Fischer. Iconicity: Literary texts. In K. Brown, editor, *Encyclopedia of Language and Linguistics*, volume 5, pages 462–472. Elsevier, Oxford, 2nd edition, 2006.
- [17] A. Nogueiras, A. Moreno, A. Bonafonte, and J. B. M. no. Speech emotion recognition using Hidden Markov Models. In *Proc. of Eurospeech*, 2001.
- [18] J. W. Pennebaker, R. J. Booth, and M. E. Francis. *Linguistic Inquiry and Word Count: LIWC*. [Computer software]. Mahwah, NJ: Erlbaum., 2001.
- [19] E. Sapir. A study in phonetic symbolism. *Journal of Experimental Psychology*, (12):225–239, 1929.
- [20] K. R. Scherer. Vocal communication of emotion: A review of research paradigms. *Speech Communication*, 40(1):227–256, 2003.
- [21] C. Strapparava, A. Valitutti, and O. Stock. The affective weight of the lexicon. In *Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC 2006)*, Genoa, Italy, 2006.
- [22] I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, San Francisco, 2nd edition edition, 2005.