Chapter 9

Validating the Coverage of Lexical Resources for Affect Analysis and Automatically Classifying New Words along Semantic Axes

Gregory Grefenstette[1]
CEA/LIST/DTSI/SCRI/LIC2M
Commissariat à l'Energie Atomique, Centre de Fontenay-aux-Roses, B.P. 6
92265 Fontenay-aux-Roses Cedex, France
Email: Gregory.Grefenstette@cea.fr

Yan Qu and David A. Evans
Clairvoyance Corporation, 5001 Baum Boulevard, Suite 700
Pittsburgh, PA 15213-1854 USA
Email: {yqu,dae,jimi}@clairvoyancecorp.com

James G. Shanahan[1]
Turn Inc
1400 Fashion Island Blvd, Suite 510
San Mateo, CA 94404, U.S.A.
Email: James Shanahan@turn.com

Abstract

In addition to factual content, many texts contain an emotional dimension. This emotive, or affect, dimension has not received a great amount of attention in computational linguistics until recently. However, now that messages (including spam) have become more prevalent than edited texts (such as newswire), recognizing this emotive dimension of written text is becoming more important. One resource needed for identifying affect in text is a lexicon of words with emotion-conveying potential. Starting from an existing affect lexicon and lexical patterns that invoke affect, we gathered a large quantity of text to measure the coverage of our existing lexicon. This chapter reports on our methods for identifying new candidate affect words and on our evaluation of our current affect lexicons. We describe how our affect lexicon can be extended based on results from these experiments.

Keywords: affect lexicon, emotion, lexicon discovery, semantic axes.

[1] This work was done while the first author was an employee of Clairvoyance Corporation.

1. Introduction

The emotive, or affective, component of text has received revived attention in computational linguistics recently. As messages (e-mail including spam, short messages, electronically submitted user opinion) become more prevalent on the Internet than edited text (such as newswire), recognizing the emotion contained in a text is becoming important as a filtering tool. All language users know that the same message content can be delivered with a wide variety of affective nuances. For example, one can present the same event as a glorious or horrible thing through judicious word choice. While the facts concerning the event may remain the same (*who, what, when, how*), different lexical selections, grammatical choices, and different focus can change the affect register of a text. Recognizing affect completely in a text would require at least recognizing rhetorical structures and emotion-bearing words. Automatic recognition of rhetorical structure is still in its infancy (Teufel and Moens, 2002), but work on the emotive content of words has a long history in linguistics[2].

1.1 Early Work on Affect Labelling

In psychological research in the early 1960s, Deese (1964) postulated that words were stored internally along semantic axes, and elaborated experiments in free association that were used to predict which words were found along axes such as “big-small”, “hot-cold”, etc. These ideas entered the field of linguistics as a “linguistic scale,” defined by Levinson (1983) as set of alternate or contrastive expressions that can be arranged on an axis by degree of semantic strength along that dimension, and also somewhat in the idea of semantic fields (Berlin and Kay, 1969; Lehrer, 1974) which correspond to a group of words that cover and divide up some semantic dimension, such as “colors.”

In addition to these lines of research interested in placing terms along semantic axes, other researchers such as Stone and Lasswell began building lexicons in which words were explicitly labeled with affect. For example, in the Lasswell Value Dictionary (Lasswell and Namenwirth, 1969), the word *admire* was tagged with a positive value along the dimension *RESPECT*. This dictionary marked words with binary values along eight basic value dimensions (WEALTH, POWER, RECTITUDE, RESPECT, ENLIGHTENMENT, SKILL, AFFECTION, and WELLBEING). Stone’s work on the General Inquirer dictionary (Stone et al., 1966) has continued to this day (see http://www.wjh.harvard.edu/~inquirer/inqdict.txt for an online version). Currently (in mid 2004) the dictionary contains 1,915 words marked as generally positive and 2,291 words marked as negative. In addition to these two general classes, a wide variety of other affect classes are used to label entries, e.g., Active, Passive, Strong, Weak, Pleasure, Pain, Feeling (other than pleasure or pain), Arousal, Virtue, Vice, Overstated, Understated. The dictionary also includes an open-ended of set of semantic labels, e.g., *Human*, *Animate*, …, *Region*, *Route*,…, *Object*, *Vehicle*,…, *Fetch*, *Stay*,… (see http://www.wjh.harvard.edu/~inquirer/homecat.htm for an explanation of these labels). In these dictionaries, all labels are binary. For example, in the General Inquirer, the word *admire* has the labels (among others) corresponding to *Positive* and *Pleasure*. Words either possess the attribute or not; there is no question of degree.

[2] Computer recognition of emotion in human faces is a possibly related and now dynamic line of research. For one reference, see Brave and Nass (2002).

1.2 Recent Work

In addition to these manually labeled lexicons of affect words, recent experiments have attempted to find labels such as *positive* and *negative* automatically via statistical corpus analysis. Hatzivassiloglou and McKeown (1997) took a number of frequently occurring adjectives that they decided had an orientation and then used statistics on whether two adjectives appeared together in a corpus in the pattern *X and Y* to automatically classify adjectives as having positive or negative orientation. Essentially, words that co-occurred with each other in that pattern were considered as having the same polarity, and the bigger class of words was considered as having negative polarity (since there are more negative words than positive words in English). They achieved 92% accuracy over a set of 236 adjectives that they classified as positive or negative. Wiebe (2000) used a seed set of "subjective" adjectives and a thesaurus generation method (Hindle, 1990) to find more subjective adjectives. Turney and Littman (2003) found another effective way of deciding whether a word can be considered as positively or negatively charged. Given a set of words that they knew to be positively or negatively charged (using tagged words from Hatzivassiloglou's and McKeown's (1997) experiments and from the General Inquirer Lexicon), they tested how often each word would appear in the context of a set of positive paradigm words (*good*, *nice*, *excellent*, *positive*, *fortunate*, *correct*, *superior*) and a set of negative paradigm words (*bad*, *nasty*, *poor*, *negative*, *unfortunate*, *wrong*, *inferior*). Using a form of point-wise mutual information (Church and Hanks, 1989) and page statistics on word appearance on Altavista and word co-occurrence (within a window of ten words using the Altavista NEAR operator[3]) they classified as positively charged words the words that appeared most significantly with the set of positive paradigm words; and as negatively charged those appearing significantly more often with the negative paradigm words. Using this method, they achieved an accuracy of 98.2% with the 334 most frequently found adjectives in the Hatzivassiloglou and McKeown test set.

1.3 Our Approach

Both groups, Turney and Littman and Hatzivassiloglou and McKeown, begin with a set of words that they consider to be emotionally charged. In our experiments as described below, we try to find words that are probably negatively or positively charged automatically, in order to extend an existing lexicon of affect words.

We see affect words as occupying a ground between stop words, e.g., *the, in, a, is* and content words, e.g., *electricity, transfer, merger*. The boundary is not clear and distinct, as sometimes affect is carried by choice of different content words, e.g., *insurgent* or *terrorist*. And what qualifies as an affect word is ultimately a subjective decision. This notwithstanding, we propose here a method for evaluating the coverage of an affect lexicon, and we demonstrate a means for extending it.

2. The Current Clairvoyance Affect Lexicon

Beginning in the late 1990s, in connection with our development of text-mining configurations of Clairvoyance technology, we began exploring the "extra-semantic" dimensions of text, including emotion. At that time we developed a lexicon of affect words by hand (Subasic and Huettner, 2000a, 2000b, 2001; Huettner and Subasic, 2000). Entries in this lexicon consist of five fields: (i) a lemmatized word form, (ii) a simplified part of speech [adjective, noun, verb, adverb], (iii) an

[3] This operator was unfortunately eliminated from Altavista in the Spring of 2004.

affect class, (iv) a weight for the centrality of that word in that class, and (v) a weight for the intensity of the word in that class. The centrality of a word is a hand-assigned value between 0.0 and 1.0 that is intended to capture the relatedness of the word to the affect class. The intensity value attempts to capture the emotional strength of the word. For example in the sample entries given below, one sees that the adjective *gleeful* has been assigned to two affect classes (*happiness* and *excitement*) and that it has been deemed more related to the class *happiness*, with a centrality of 0.7, than it is to the class *excitement*, where the lexicon creators only gave it a centrality of 0.3.

"gleeful"	adj	happiness	0.7	0.6
"gleeful"	adj	excitement	0.3	0.6

In both entries, the word *gleeful* was deemed to have an intensity of 0.6 (out of a maximum intensity of 1). The combination of intensities and centralities made it possible to develop multidimensional weightings of affect in texts (Subasic and Huettner, 2000a, 2000b, 2001).

The existing lexicon contains 3,772 entries. A word form, such as *gleeful*, can appear in more than one entry. There are 2,258 different word forms (ranging from *abhor, abhorrence, abject, absurd, abuse, abusive, acclaim, accomplish* to *worth, wrong, wrongdoing, yawn, yearn, yearning, yen, yucky*). There are 86 different affect classes[4], such as *happiness* and *excitement* shown above. The numbers of entries for each affect class are given in Table 1.

Positive class	*Negative class*		*Positive class*	*Negative class*
Advantage (46)	Disadvantage (59)		Love (37)	Hate (28)
Amity (26)	Anger (28)		Loyalty (20)	Disloyalty (19)
Attraction (71)	Repulsion (70)		Morality (25)	Immorality (64)
Clarity (18)	Confusion (56)		Nurturance (35)	Harm (108)
Comfort (0)	Irritation (56)		Openness (47)	Slyness (82)
Cooperation (21)	Conflict (141)		Peace (21)	Violence (139)
Courage (44)	Fear (71)		Persuasion (26)	Force (102)
Creation (42)	Destruction (75)		Pleasure (33)	Pain (69)
Desire (39)	Avoidance (72)		Praise (36)	Slander (59)
Energy (27)	Fatigue (42)		Predictability (37)	Surprise (56)
Excitement (77)	Boredom (30)		Promise (22)	Warning (24)
Facilitation (13)	Prevention (36)		Public-spiritedness (1)	Crime (77)
Happiness (23)	Sadness (40)		Reasonableness (27)	Absurdity (19)
Health (13)	Sickness (14)		Responsibility (21)	Irresponsibility (33)
Honesty (21)	Deception (83)		Sanity (16)	Insanity (42)
Humility (24)	Pride (46)		Security (25)	Insecurity (24)
Humor (23)	Horror (50)		Selflessness (25)	Greed (55)
Innocence (20)	Guilt (40)		Sensitivity (16)	Insensitivity (32)
Intelligence (49)	Stupidity (32)		Strength (42)	Weakness (57)
Justice (70)	Injustice (41)		Success (43)	Failure (54)
Lively (0)	Death (31)		Superiority (108)	Inferiority (67)
			Surfeit (45)	Lack (75)

Table 1. List of paired (positive-negative) affect classes in the existing Clairvoyance lexicon, with number of headwords present for each class.

[4] The Humanity Quest Web site lists more than 500 different human values, similar to our affect classes. See http://web.archive.org/web/20031118174947/http://humanityquest.com/.

In any practical text-analysis application the question always arises whether the lexical resources are sufficient. In our case, we are interested in knowing whether our affect lexicon is complete. To answer this question, we decided to mine the Web using lexical patterns that we thought might be productive indicators of affect words. These patterns are described in the next section.

3. Emotive Patterns

Insults are highly charged with emotional content. Typical insults might be: "he is such a jerk/idiot/know-it-all!" The same pattern "he is such a …" can also introduce a complimentary characterization: "he is such a prince/magnificent artist/all-around player!" After exploring a few such patterns by typing them into a Web browser and seeing what was brought back, we decided to test the patterns generated by the following procedure systematically:

> Create a pattern by constructing a two word phrase composed of one of these 21 words:
> > {*appear, appears, appeared, appearing, feel, feels, feeling, felt, are, be, is, was, were, look, looked, looks, looking, seem, seems, seemed, seeming*}
>
> followed by one of the 5 words:
> > {*almost, extremely, so, too, very*}

For each of these 105 patterns, e.g., "looking extremely…", we sent off a search request and extracted up to 4,000 text snippets containing the pattern from the results pages on www.alltheweb.com[5]. From each context snippet, we extracted the word appearing directly after the pattern. For example, for the pattern "looking extremely," we extracted "dubious" from the following snippet:

> The Christian Science Monitor: Hands-on art gets a grip on athletes inner self
> Famed baseball star Sammy Sosa is standing in a conference room in a downtown hotel here, **looking extremely dubious** about placing his hand in a pan of hot wax. Sculptor Raelee Frazier (in photos at right with Sosa) guides his right ...

The most common words appearing after this particular pattern "looking extremely" were the following:

77 good
52 pleased
47 uncomfortable
41 bored
40 happy
38 promising
35 tired
27 pissed
27 pale

For example, "looking extremely promising" appeared in 38 of the 4000 snippets.

When we produce similar statistics for all the words appearing after any of the 105 patterns, we get the following list:

[5] 4,000 was the maximum number of page results that one could obtain from the AllTheWeb browser in 2004. Google and AltaVista limited their responses to 1,000 pages.

8957 good
2906 important
2506 happy
2455 small
2024 bad
1976 easy
1951 far
1745 difficult
1697 hard
1563 pleased

There were 15,111 different, inflected words found at least once immediately following the 105 patterns on the results pages returned by AllTheWeb. Although these patterns seem to give many affect words, e.g., *good, bad*, not all words, even at the top of the list, are affect words. In the next sections, we describe how we can judge whether a pattern is productive for finding affect words.

3.1 Two Gold Standards for Identifying Affect Words

In order to measure the productivity of these patterns, one of the authors examined each of the 4,746 words that appeared more than twice (out of the 15,111 words found at least once) after the patterns and decided subjectively, without referring to the existing affect lexicon, whether the word should be considered an emotion-bearing, affect word (2,988 words) or not (1,758 words). Some of the most frequently appearing words that were marked as an affect word by this author were: *good, important, happy, bad, easy, difficult, hard, pleased, nice, proud, comfortable, tired, helpful, impossible, busy*. Some of the most frequently appearing words that the author did not consider to be affect words were: *small, far, similar, different, high, long, large, close, simple, big, identical, exactly, low, late, real*. We call this adjudicated list the Manual Gold Standard (MGS) in our evaluations.

A second gold standard was produced by listing all the words found in the General Inquirer Lexicon that possessed one of the following affect-related labels: *Pos, Neg, Pstv, Ngtv, Negate, Hostile, Strng, Power, Weak, Subm, Pleasure, Pain, Arousal, EMOT, Feel, Virtue, Vice, IAV, SV, IPadj, IndAdj, EVAL*. (See http://www.wjh.harvard.edu/~inquirer/homecat.htm for explanation of these categories.) Of the 9,051 different headwords found in the General Inquirer Lexicon, 5,574 possessed at least one of these labels, and 3,477 others did not. We will call this set the General Inquirer Gold Standard (GIGS).

Given these two gold standards of affect/non-affect words, we judged both the productivity of the emotive patterns, as well as the coverage of our existing affect lexicon.

3.2 Evaluating the Productivity of Emotive Patterns

Each emotive pattern, e.g., "appears almost…", was evaluated by referring to the gold standard lists of affect/non-affect words described in the previous section. We tabulated the number of words produced by the pattern and found in the gold standard (in the column labeled *found*), the number of these words that the gold standard had listed as an affect word (= *good*) or non-affect word (= *bad*). If a candidate word found by the pattern was not in the gold standard, we did not count it. We discuss these cases below.

Table 2 shows the results for each set of emotive patterns against the Manual Gold Standard. Not all 105 patterns are shown here; each line corresponds to all the variants of a word given in the first column. For example, all the results for “seem almost”, “seems almost”, “seemed almost” and “seeming almost” are collated in the first line. In that line, we find that these patterns picked up 1,254 of the 4,746 words found in the Manual Gold Standard of affect/non-affect words, with 957 of these 1,254 words (precision 76%) corresponding to affect words. The patterns involving “extremely” had the best precision. and the patterns consisting of versions of “be so” had the best recall of words from the Manual Gold Standard, picking out 1,465 of the 2,988 affect words found there, but with a precision of only 71%.

Emotive pattern		*found*	*good*	*bad*	*precision*
seem	almost	1254	957	297	0.76
seem	extremely	1170	973	197	0.83
seem	so	1372	1095	277	0.80
seem	too	1220	1006	214	0.82
seem	very	1216	977	239	0.80
feel	almost	1092	785	307	0.72
feel	extremely	1082	860	222	0.79
feel	so	1086	830	256	0.76
feel	too	1120	844	276	0.75
feel	very	1160	905	255	0.78
appear	almost	1063	647	416	0.61
appear	extremely	*618*	*518*	**100**	**0.84**
appear	so	1170	857	313	0.73
appear	too	1178	897	281	0.76
appear	very	1344	1041	303	0.77
look	almost	996	667	329	0.67
look	extremely	1305	1014	291	0.78
look	so	1055	798	257	0.76
look	too	1106	755	351	0.68
look	very	1066	801	265	0.75
be	almost	1320	680	*640*	*0.52*
be	extremely	1541	1157	384	0.75
be	so	**2053**	**1465**	588	0.71
be	too	1393	1019	374	0.73

Table 2. Productivity and precision of emotive patterns against the Manual Gold Standard. The data in each row are summed over all variants of the word form in the first column, e.g., “appear too” covers results from “appears too”, “appearing too”, etc. The third column shows how many words in the gold standard were found after the patterns. The fourth and fifth columns show how many of these discovered words were marked as Affect or Non-Affect, respectively, in the gold standard. The last column shows the precision of the pattern for finding Affect words. The best numbers are shown in bold and the worst are shown in italics.

Table 3 shows the results for sets of emotive patterns against the General Inquirer Gold Standard. The patterns involving “extremely” once again had the best precision, with the patterns with “so” a close second. As before, the patterns consisting of versions of “be so” had the best recall of words from the gold standard, picking out 1,026 of the 5,574 General Inquirer lexicon words possessing an affect label, with a precision of 83%.

These results demonstrate that it is possible to identify lexical patterns for finding emotion-bearing, affect words with a high precision. The patterns that we used establish contexts for (and, hence, find) adjectives and participles. Other patterns must be used to find verbs and nouns, e.g., maybe a pattern such as *never dare to X* to select verbs.

Emotive pattern		*found*	*good*	*bad*	*precision*
seem	almost	743	584	159	0.79
seem	extremely	760	656	104	0.86
seem	so	902	766	136	0.85
seem	too	797	664	133	0.83
seem	very	828	693	135	0.84
feel	almost	564	439	125	0.78
feel	extremely	*547*	478	**69**	**0.87**
feel	so	600	512	88	0.85
feel	too	588	477	111	0.81
feel	very	657	547	110	0.83
appear	almost	630	447	183	0.71
appear	extremely	448	380	68	0.85
appear	so	811	685	126	0.84
appear	too	745	606	139	0.81
appear	very	911	766	145	0.84
look	almost	581	*426*	155	0.73
look	extremely	771	634	137	0.82
look	so	691	592	99	0.86
look	too	698	550	148	0.79
look	very	679	564	115	0.83
be	almost	843	576	*267*	*0.68*
be	extremely	1055	865	190	0.82
be	so	**1240**	**1026**	214	0.83
be	too	920	726	194	0.79
be	very	1163	924	239	0.79

Table 3. Productivity and precision of emotive patterns against the General Inquirer Gold Standard. The data in each row are summed over all variants of the word form in the first column, e.g., "appear too" covers results from "appears too", "appearing too", etc. The third column shows how many words in the gold standard were found after the patterns. The fourth and fifth columns show how many of these found words were marked as Affect or Non-Affect, respectively, in the gold standard. The last column shows the precision of the pattern for finding Affect words. The best numbers are shown in bold and the worst are shown in italics.

3.3 Evaluating the Coverage of Existing Affect Lexicons

The second part of our evaluation concerns verifying how many of the affect words identified as such in the hand-tagged gold standard are actually found in the existing affect lexicon developed in previous work (Subasic and Huettner, 2000a, 200b, 2001). There are 2,988 affect words marked in our Manual Gold Standard. Of these words, only 655 were found in our existing lexicon, which therefore has a coverage of 22%. Some of the words from the emotive patterns tested that were not in our lexicon are: *difficult, pleased, nice, comfortable, impossible, busy, young, old, strongly, hot, uncomfortable, expensive, interested, strange, interesting, lucky, sorry,* *normal, cold, familiar, grateful, professional, new, natural, complex, pretty, welcome, light, relaxed, rare, fast, likely, special, limited, early, lonely, serious, tight, vulnerable, certainly, upset, sweet, blessed, positive, human, unfashionable, unflattering, ungrounded, unhelpful, unhip, unimaginably, uninhibited, unintelligent, uninvolved, unladylike, unmanageable, unmatched, unnerving, unnoticeable, unpalatable, unpolished, unproductive, unqualified, unquestionable, unread, unresponsive, unrestricted, unruly, unsatisfying, unsexy, unspecific, unsuitable, unwatchable, unwelcoming, upfront, uppity, venomous, victimized, vindicated, virile, visceral, wan, watchful, weighty, weirded, wellcome, well-kept, well-qualified, well-read, well-researched.*

The intersection between the General Inquirer and our existing Clairvoyance affect lexicon contains only 1,292 of the 5,574 affect tagged words. Some of these missing words are *abandon, abandonment, abate, abdicate, abide, able, abnormal, abolish, abominable, abound, abrasive, abrupt, abscond, absence, absent, absent-minded, absentee, absolute, absolve, absorbent, absorption, absurdity, abundance, abundant, abyss, accede, accelerate, acceleration, accentuate, accept, acceptable, acceptance, accessible, accession, accident, acclamation, accolade, accommodate, accommodation, accompaniment,…*

The intersection between the Manual Gold Standard and the General Inquirer Gold Standard has 1,295 words. Here are some words not marked with affect labels listed above in the General Inquirer: *young, impressed, slow, complicated, relaxed, obvious, likely, concerned, early, tight, embarrassed, dry, knowledgeable, exclusively, totally, sexy, inclined, instantly, informative, distant, overwhelmed, quickly, quick, nonexistent, carefully, effortless, crowded, isolated, surreal, exhausted, personal, finished, stressed, detailed, easily, sleepy, diverse, loose, restrictive, annoyed,…* Some of these words can be derived from words appearing in the General Inquirer, e.g., *relaxed* from *relax, annoyed* from *annoy,* and other words are not marked with affect labels but with other labels (e.g., *young* is marked as a time interval) since the purpose of the General Inquirer Lexicon is not limited to analyzing affect, but to serve as a resource for more general text analysis.

From this analysis it can be concluded that the definitive affect lexicon has not yet been created and that there is room for improvement in existing affect lexicons.

4. Scoring the Intensity of Candidate Affect Words

In the previous section, it was stated that the 105 emotive patterns had uncovered 4,746 words that appeared in the patterns three or more times from the snippets retrieved. These words were classified by hand as affect or non-affect bearing words to form our Manual Gold Standard. Rather than classifying these terms by hand, as we did to create the gold standard, one might use the automatic ranking technique for calculating the polarity of an unknown word described by Turney and Littman (2003). We replicated this technique of calculating point-wise mutual information with positive-negative paradigm words described in this article and applied it to the words extracted by our emotive patterns.

Our application of this technique proceeded as follows. Each candidate affect word was used to create a series of 14 requests to AltaVista. Each request placed the word with one of the paradigm words using the NEAR operator. For example, given the word *comfortable*, a series of AltaVista requests was created as shown in Table 4, in which we also show the number of pages that AltaVista found containing the pair of words near each (within ten words, according to AltaVista).

Negative Paradigm Queries	*Page Counts*
comfortable NEAR bad	13127
comfortable NEAR nasty	1008
comfortable NEAR poor	6943
comfortable NEAR negative	3836
comfortable NEAR unfortunate	535
comfortable NEAR wrong	8449
comfortable NEAR inferior	437
Positive Paradigm Queries	*Page Counts*
comfortable NEAR good	184024
comfortable NEAR nice	57757
comfortable NEAR excellent	95119
comfortable NEAR positive	13259
comfortable NEAR fortunate	1276
comfortable NEAR correct	7952
comfortable NEAR superior	39182

Table 4. Example of some of the raw data used in the Turney and Littman (2003) method. For a given word, here "comfortable," one sends requests to a Web search engine to find how many times the word co-occurs with negative connotation words or with positive connotation words. Here we show the page counts from Altavista.

Nwords	*Pages*	*Pwords*	*Pages*
bad	24576337	good	54596054
nasty	3712598	nice	17084308
poor	10813343	excellent	15955669
negative	7430078	positive	11797788
unfortunate	1174016	fortunate	1357375
wrong	13037886	correct	14187506
inferior	1565672	superior	9377519

Table 5. To calculate point-wise mutual information, one also needs the page counts of the negative and positive paradigm words, such as given here by Altavista in early 2004.

$$\text{SO - PMI}(word) = \log_2\left(\frac{\prod_{pword \in pwords} \text{hits}(word \text{ NEAR } pword) \bullet \prod_{nword \in Nwords} \text{hits}(nword)}{\prod_{pword \in pwords} \text{hits}(pword) \bullet \prod_{nword \in Nwords} hits(word \text{ NEAR } pword)}\right)$$

Figure 1. The point-wise mutual information formula from Turney and Littman (2003). "Pwords" is the set of positive paradigm words (here, as in that article, we used the set {good, nice, excellent, positive, fortunate, correct, superior}) and "Nwords" is the set of negative paradigm words ({bad, nasty, poor, negative, unfortunate, wrong, inferior}).

Using the point-wise mutual information formula from Turney and Littman (2003), shown in Figure 1, and the Altavista page statistics for the positive paradigm *Pwords* and the negative paradigm *Nwords*, shown in Tables 4 and 5, one finds a point-wise mutual information score for *comfortable* of 10.6553. This shows the word is more strongly associated with the positive paradigm words than the negative paradigm words, and thus is probably more positively charged. When the same calculations are applied to all the 4,746 words discovered by the affect patterns of Section 3, eliminating the 156 words that appears fewer than 100 times with all the positive and negative paradigm words (e.g., *sticklike, identical, featurerich*[6], *easytouse, shuai, goodthe*,…), we find the highest and lowest ranking words to be as given in Table 6.

SO-PMI Score	
37.5	knowlegeable
33.6	tailormade
32.9	eyecatching
29.0	huggable
26.2	surefooted
24.6	timesaving
22.9	personable
21.9	welldone
21.0	handdrawn
20.8	commonsensical
20.0	homelike
20.0	hightech
…	…
-15.1	unaccomplished
-15.6	inelegant
-15.9	spindly
-15.9	childishly
-16.2	simpleminded
-16.2	blasphemous
-17.1	underdressed
-17.5	uncreative
-18.1	disapproving
-18.5	meanspirited
-18.6	unwatchable
-22.7	discombobulated

Table 6. Point-wise mutual information scores for some words discovered by the affect patterns.

While the words with very high or very low scores seem to be affect-laden words, as Turney and Littman (2003) have found, the words around 0.0 are less clear-cut. For example, between SO-PMI scores of 0.5 and -0.5 we find word like: *jaunty, julia, jumping, kick, km, know, knowing, labor, ladies, laid, late, learned, lend, liberal, life, lit, lithe, localized, loveable, luscious, magic, main, manmade, materially, military, mindboggling, miss, missed, misty,* some of which we would classify as affect words. This SO-PMI could thus be used to rank words for inclusion in an affect dictionary, with words at extreme points (involving a threshold) included automatically and others treated manually.

[6] One step in our text processing removed hyphens from words, so a term like *feature-rich* was treated as the string *featurerich*, which leads to its low counts. *Feature-rich* (with the hyphen) appears often on the web, but was not tested in the experiments described here.

4.1 Automatically Placing Words along other Semantic Axes

Another point to consider in the case of an affect lexicon including not just positive and negative orientation but affect classes as seen in the Clairvoyance affect lexicon (see Section 2) is how to decide in which class the new words should be included. We have been experimenting with extending Turney's and Littman's (2003) technique, as they suggested, to different semantic scales. For each of the 86 affect classes (cf. Table 1) defined in our lexicon, we manually selected 4 to 6 paradigm words. For example, here are the paradigm words we chose for some of these classes:

- **praise** – *acclaim, praise, congratulations, homage, approval*
- **slander** – *bad-mouth, calumniate, calumny, defamation, slander*
- **comfort** – *comfort, comfortable, solace, comforting*
- **irritation** – *aggravate, aggravation, irritation, irritate, bothersome*
- **pleasure** – *pleasure, enjoy, delight, joy, pleasing*
- **pain –** *agony, hurt, pain, nuisance, painful, hurting*
- **excitement** – *agitate, agitation, exciting, excitement, stimulating*
- **boredom** – *boredom, boring, wearisome, tedious, tiresome*
- **humility** – hum*ble, humility, abject, modest, modesty*
- **pride** – *arrogance, arrogant, proud, prideful, pride*
- **confusion** – *confuse, confusion, unclear, confused, jumbled*

Having chosen such sets for each of the 86 affect classes, we ran queries placing the words that we wish to classify in a semantic class with the NEAR operator and each of these paradigm words as described in Section 3 with positive and negative paradigm words. Since we are not comparing polar endpoints, we used the following formula to produce a score:

$$\text{Score}(\textit{word}, \textit{Class}) = \log_2 \left(\frac{\prod_{\textit{cword} \in \textit{Class}} \text{hits}(\textit{word} \text{ NEAR } \textit{cword})}{\prod_{\textit{cword} \in \textit{Class}} \log_2(\text{hits}(\textit{cword}))} \right)$$

where *cword* is one of the paradigmatic words chosen for an affect class *Class*, and *hits* is the number of pages found by Altavista. The extra logarithm in the denominator was added only to scale the resulting score, since the counts of the *cword*s were many orders of magnitude larger than the counts of the paired words. Given a candidate affect word, this score is calculated for each of the 86 classes.

As an example of the application of this score, the word *discombobulated* scores highest with paradigm words for semantic classes: **confusion** (score: 0.634), **surprise** (-4.05), **pleasure** (-6.89),… Table 7 gives some other words and the highest scoring classes.

It would seem that this extension of Turney and Littman's (2003) technique to other semantic classes, which relies on finding co-occurrences with paradigms to choose semantic classes, similar to Turney's (2001) work on the finding of the closest synonyms, might allow us to automate the assignment of affect class centrality.

aversion	Hate (15.98), Pain (15.02), …
award	Success (33.21), Praise (33.12), …
awful	Pain(28.88), Horror (26.24), …
back-biting	Deception(4.77), Slander (4.46), …
back-stabbing	Conflict (4.16), Disloyalty (3.79), …
banditry	Security (5.48), Violence (3.49), …
barbaric	Violence (13.89), Horror (12.87), …
barbarity	Horror (9.20), Surprise (5.66), …

Table 7. Highest scoring semantic classes for some words discovered by the affect patterns, using an extension of the Turney & Littman (2003) technique to other semantic axes.

Once a class is assigned, intensity might be represented by re-using the SO-PMI formula, but with the paradigm words of the positive and negative class members. For example, for the table above, we see that *awful* is central to *Pain.* In table 1, we see that the *Pain* class is associated with the *Pleasure* class. Using the *Pleasure* paradigm words (*pleasure, enjoy, delight, joy, pleasing)* as the *Pwords* and the *Pain* paradigm words (*agony, hurt, pain, nuisance, painful, hurting*) as the *Nwords* gives us an SO-PMI $_{\text{Pleasure-Pain}}$ score for *awful* of -1.25, showing that *awful* has a moderate intensity along the negative axis *Pain*.

5. Future Work

In this chapter, we described our use of a few emotive patterns for discovering adjectives. One direction of our future work is to expand the set of emotive patterns used for extracting affect words. We have begun to mine these patterns automatically from the Web using affect words as seeds, gathering a large number of web pages containing both an existing affect word and its affect class name. From such pages, we have begun extracting patterns, e.g., *to their, the most*, *full of,* appearing before known affect words. These discovered affect patterns may yield new sets of candidate words.

Another direction for future work is confirming that the promising techniques proposed by Turney and Littman (2003) for finding negative–positive polarity can be used for automatically assigning class centrality and intensity, as the results from the previous section promise.

Another avenue to explore concerns alternative methods for placing words along a class axis. Horn (1969) proposed using the pattern *X even Y*, e.g., *silly even ridiculous*, to distinguish which element of *X* and *Y* is more intense along a scaled dimension, such as one of our affect dimensions of Table 1. Such patterns may be usefully explored on the Web as another way to align words along an axis, as suggested by Hatzivassiloglou and McKeown (1993). A quick check on Google in early 2004 shows that the contiguous phrase *silly even ridiculous* can be found on 13 pages, while *ridiculous even silly* is only found once[7]. It might be useful to try this, or other such patterns, to verify an internal ranking along a dimension.

[7] Using *silly even…* as a search patterns brings up other adjectives such *useful* and *offensive* which might not be regarded as belonging to the same affect dimension, so one stills requires a mechanism to find words in the same dimension in order to exploit this *X even Y* pattern.

6. Conclusions

We have explored a method for identifying rich sources for discovering new emotion-laden affect words via emotive lexical patterns. Using the patterns to mine the Web, we retrieved large numbers of affect words. These new words can be used to identify missing items in existing lexicons. We have shown that Turney and Littman's (2003) paradigm word co-occurrence scoring can be used to identify a certain number of the missing items as likely affect words. We have also shown that a similar technique of word co-occurrence with paradigm words might identify the likely centrality of new words among 43 pairs of positively and negatively oriented affect classes. Finally, we have preliminary results that suggest that extending Turney and Littman's approach to these semantic axes may provide an automatic way to find the intensity of new words.

7. Bibliography

Brave, S. and Nass, C. (2002) Emotion in human-computer interaction. In Jacko, J. and Sears, A. (Eds.) *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*. Lawrence Erlbaum Associates, Inc., Mahwah, NJ.

Berlin, B. and Kay, P. (1969) *Basic color terms: their universality and evolution*. University of California Press, Berkeley.

Church, K. W. and Hanks, P. (1989) Word association norms, mutual information and lexicography. In *Proceedings of the 27th Annual Conference of the Association of Computational Linguistics*. 76−82.

Deese, J. (1964) The Associative Structure of some Common English Adjectives. *Journal of Verbal Learning and Verbal Behavior 3(5)*. 347−357.

Hatzivassiloglou, V. and McKeown, K. R. (1993) Towards the automatic identification of adjectival scales: Clustering adjectives according to meaning. In *Proceedings of 31st Annual Meeting of the Association for Computational Linguistics*. 172−182.

Hatzivassiloglou, V. and McKeown, K. R. (1997) Predicting the semantic orientation of adjectives. In *Proceedings 35th Annual Meeting of the Association for Computational Linguistics*. 174−181.

Hindle, D. (1990) Noun classification from predicate argument structures. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*. 268−275.

Horn, L. (1969) A Presuppositional Analysis of Only and Even. In *Papers from the 5th Regional Meeting of the Chicago Linguistics Society*. 98−107.

Huettner, A. and Subasic, P. (2000) Fuzzy Typing for Document Management. In *ACL 2000 Software Demonstration*.

Lasswell, H. D. and Namenwirth, J. Z. (1969) *The Lasswell Value Dictionary*. Yale University Press, New Haven.

Lehrer, A. (1974) *Semantic Fields and Lexical Structure*. North Holland, London.

Levinson, S. C. (1983) *Pragmatics*. Cambridge University Press, Cambridge, UK.

Stone, P. J., Dunphy, D. C., Smith, M. S., and Ogilvie, D. M. (1966) *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge, MA.

Subasic, P. and Huettner, A. (2000a) Affect Analysis of Text Using Fuzzy Semantic Typing. In *Proceedings of FUZZ-IEEE 2000*.

Subasic, P. and Huettner, A. (2000b) Calculus of Fuzzy Semantic Typing for Qualitative Analysis of Text. In *Proceedings of ACM KDD 2000 Workshop on Text Mining*.

Subasic, P. and Huettner, A. (2001) Affect Analysis of Text Using Fuzzy Semantic Typing. *IEEE Transactions on Fuzzy Systems, Special Issue*.

Teufel, S. and Moens, M. (2002) Summarizing Scientific Articles – Experiments with Relevance and Rhetorical Status*. Computational Linguistics, 28 (4)*.

Turney, P. D. (2001) Mining the Web for Synonyms: PMI-IR versus LSA on TOEFL. In *Proceedings of the Twelfth European Conference on Machine Learning (ECML2001)*. 491−502.

Turney, P. D. and Littman, M. L. (2003*)* Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS), 21(4)*, 315−346.

Wiebe, J. (2000) Learning subjective adjectives from corpora. In *Proceedings of AAAI 2000*. 735−740.