

Chapter 1

Contextual Valence Shifters

Livia Polanyi

FXPAL

3400 Hillview Ave, Bldg. 4

Palo Alto CA 94304

E-mail: polanyi@fxpal.com

Annie Zaenen

PARC,

3333 Coyote Hill Road,

Palo Alto, CA 94304 USA

E-mail: zaenen@parc.com

Abstract

In addition to describing facts and events, texts often communicate information about the attitude of the writer or various participants towards material being described. The most salient clues about attitude are provided by the lexical choice of the writer but, as discussed below, the organization of the text also contributes information relevant to assessing attitude. We argue that the current work in this area that concentrates mainly on the negative or positive attitude communicated by individual terms (Edmonds and Hirst, 2002; Hatzivassiloglou and McKeown, 1997; Turney and Littman, 2002; Wiebe et al., 2001) is incomplete and often gives the wrong results when implemented directly. We then describe how the base attitudinal valence of a lexical item is modified by lexical and discourse context and propose a simple, “proof of concept” implementation for some contextual shifters.

Keywords: attitude, discourse, valence shifters, genre structure, multiple constraints, calculating valence.

1. Introduction

In addition to describing facts and events, texts often communicate information about the attitude of the writer or various participants towards an event being described. Salient clues about attitude are provided by the lexical choice of the writer but, as discussed below, the organization of the text also contributes critical information for attitude assessment. We start from the current work in this area that concentrates mainly on the negative or positive attitude communicated by individual

terms (Edmonds and Hirst, 2002; Hatzivassiloglou and McKeown, 1997; Turney and Littman, 2002; Wiebe et al. 2001). We argue that this approach is incomplete and often gives the wrong results when implemented directly. We describe how the base attitudinal valence of a lexical item can be modified by context and propose a simple “proof of concept” implementation for some contextual shifters.

2. From Simple Valence to Contextually Determined Valence

2.1 Simple Lexical Valence

Examples of lexical items that communicate a negative or positive attitude (*valence*) can be found in all open word classes and as multi-word collocations such as *freedom fighter*. Below we have listed some examples of English words which can be readily characterized as positively or negatively valenced¹.

PART OF SPEECH	Positive Valence	Negative Valence
Verbs	Boost, Embrace, Ensure, Encourage, Delight, Manage, Ease	Conspire, Meddle, Discourage, Fiddle, Fail, Haggle
Nouns	Approval, Benefit, Chance, Approval, Benefit, Credit, Favor, Freedom, Hope	Backlash, Backlog, Bankruptcy, Beating, Catastrophe
Adjectives	Attractive, Better, Brave, Bright, Creative, Dynamic,	Annoying, Awry, Arbitrary, Bad, Botched,
Adverbs	Attractively, ...	Annoyingly, ...

Table 1. Examples of words with non-neutral valence.

2.2 Lexical Valence in Texts

To illustrate how lexical valence influences interpretation, let us look briefly at three short texts.² While all of the texts communicate the same denotative information, the connotative force of each version is different. In the first text, the protagonist is an unremarkable young man, in the second text, he is a much friendlier, warmer sort of chap while he emerges in the third text as a juvenile delinquent³:

Text 1. The eighteen year old walked through the part of town where he lived. He stopped for a while to talk with people on the street and then went to a store for some food to bring to the small apartment where he lived with some people he knew.

¹ Not all terms can be characterized along this dimension: many terms are essentially neutral.

² Space constraints and the difficulty of finding short texts that exemplify important complex cases while presenting few other distractions oblige us to construct our example.

³ Notation: Relevant terms are bold; positive terms are marked with a +; negative terms are marked with a -; comparable neutral terms are underlined.

Text 2. *The young man⁺ strolled⁺ through his neighborhood⁺. He lingered⁺ to chat⁺ with people on the street and then dropped into⁺ a shop⁺ for some goodies⁺ to bring home⁺ to the cozy⁺ place which he shared⁺ with some friends⁺.*

Text 3. *The teenaged male⁻ strutted⁻ through his turf. He loitered⁻ to shoot the bull⁻ with people on the street and then ducked⁻ into a dive⁻ for some grub⁻ to bring to the cramped hole-in-the-wall⁻ where he crashed⁻ with some cronies⁻.*

The difference in perlocutionary force among these texts emerges solely from the combined effects of the choice of synonyms (or near synonyms) chosen to depict the persons, events and situation involved.

Observations such as these have led researchers to classify terms as positive or negative. The simple computation of the attitude expressed in a text would then consist of counting the negative and positive instances and decide on the basis of the highest number.

To see that the simple counting will not work for many texts, consider the following example (from *The Economist*):

Of course, that would not stop deregulation of the power industry altogether. The blunderbuss⁻ of state initiatives will see to that. However, by prolonging uncertainty⁻, it would needlessly delay⁻ the arrival of the bonanza⁺ of benefits⁺ that consumers deserve⁺, and give them legitimate⁺ grounds for their cynicism⁻.

While there are six negative lexical items (marked with -) and only four positive items (marked with +) in this text, readers do not conclude that the author is negative about “deregulation”. In fact, the writer views deregulation positively. Clearly, then, the full story of how lexical items reflect attitudes is more complex than simply counting the valences of terms would suggest. In the remainder of this paper, we will propose a number of ways in which the basic valence of individual lexical items may be strengthened or weakened by context provided by (1) the presence of other lexical items, (2) the genre type and discourse structure of the text and (3) cultural factors.

In looking at texts, it is clear that lexical items can be strongly positive or negative or somewhat strong or weak or “hint” at a positive or negative connotation. Therefore, characterizing terms in binary terms as either positive or negative as we have done so far is too crude. Believing that it would be desirable to have a more fine-grained classification, we have adopted a slightly more sensitive scale with three positive and three negative values. In the notation we adopt in this paper, therefore, we assume that words like *clever* and *successful* are marked +2 in the lexicon. Negatively valenced items are marked -2. It should be kept in mind, however, that this scheme falls far short of an adequate solution to this problem.

3. Contextual Valence Shifters

3.1 Sentence Based Contextual Valence Shifters

While some terms in a text may seem to be inherently positive or negative, we shall show how others change base valence according to context – receiving their perlocutionary force either from the domain of discourse or from other lexical items nearby in the document. In the remainder of this paper we will discuss a number of interacting factors that make the determination of the point of view that an author expresses in a document difficult. We will begin with a survey of several lexical phenomena that can cause the valence of a lexical item to shift from one pole to the other or, less forcefully, to modify the valence towards a more neutral position.

3.1.1 Negatives and Intensifiers

The most obvious shifters are negatives.⁴ How *not* can flip the valence of a term has been discussed in the computational literature (Das and Chen, 2001; Pang, Lee and Vaithyanathan 2002). However, in addition to *not*, negatives can belong to various word classes. Simple negatives include *never*, *none*, *nobody*, *nowhere*, *nothing*, and *neither*. For example:

John is clever versus *John is not clever*.

John is successful at tennis versus *John is never successful at tennis*.

Each of them is successful versus *None of them is successful*.

Combining positively valenced words with a negation such as *not* flips the positive valence to a negative valence. For example⁵:

clever +2 combined with *not* \Rightarrow *not clever* -2

successful +2 combined with *not* \Rightarrow *not successful* -2

The combination of a positive evaluator with a negation turns the evaluation as a whole into a negative one. Inversely the combination of a negative evaluator with negation turns the whole into a positive evaluation (e.g., “*He is not stupid*”).

Not all modifiers switch the valence. Intensifiers such as the *rather* in *rather efficient* and the *deeply* in *deeply suspicious* act to weaken or strengthen the base valence of the term modified. *Rather* weakens the force of a term and *deeply* enhances it (Riloff and Wiebe, 2003). We can calculate their effect by adding or subtracting a ‘point’ to/from the base valence of a term.

Suspicious -2 \Rightarrow *deeply suspicious* -3

Efficient +2 \Rightarrow *rather efficient* +1

As with the negative shifters, intensifiers can belong to all open lexical classes. In addition to adverbs, quantifiers such as *few*, *most*, and nouns such as *lack (of)* also exist.

3.1.2 Modals

Language makes a distinction between events or situations which are asserted to have happened, are happening or will happen (*realis* events) and those which *might*, *could*, *should*, *ought to*, or *possibly* occurred or will occur (*irrealis* events). **Modal operators** set up a context of possibility or necessity and in texts they initiate a context in which valenced terms express an attitude towards entities which do not necessarily reflect the author’s attitude towards those entities in an actual situation under discussion. Therefore, in computing an evaluation of the author’s attitude, terms in a modal context should not be treated precisely as terms in a *realis* context.

Assume the *realis* sentences: *Mary is a terrible person. She is mean to her dogs. Terrible* and *mean* are negatively valenced terms. The score for each of the sentences is -1. However, the sentence *If Mary were a terrible person, she would be mean to her dogs*, asserts neither that *Mary*

⁴ Of course for a shift in attitude to take place there has to be an attitude expressed in the first place. A simple sentence such as “*John is home*” might express a simple fact without betraying an attitude (i.e. the attitude score is 0). When negated, as in “*John is not home*”, there is no shift in attitude (i.e. the negation of 0 is 0).

⁵ While it is a simplification to take the scope of a negative as always a whole clause, we will assume this here.

is a terrible person nor *that she is mean to her dogs*. On the contrary, the force of *would* suggests that *she is not mean to her dogs* while the *If* sets up a context in which *Mary is not necessarily a terrible person*. In fact, we tend to believe that she is not *terrible* at all. Therefore, the modal operators neutralize the base valence of *terrible* and *mean*, resulting in a re-computed value of 0 for the modal version.

3.1.3 Presuppositional Items

Often words shift the valence of evaluative terms through their presuppositions. This is typical for adverbs like *barely* as shown by comparing “It is sufficient” with “It is *barely sufficient*.” “Sufficient” is a positive term, “*barely sufficient*” is not: it presupposes that better was expected. These terms can introduce a negative or a positive evaluation even when there are no other evaluative terms around, as in *He got into Foothill College* versus *He barely got into Foothill College* or *He got into Harvard* and *He even got into Harvard*. Words like *barely* and *even* will be marked in the lexicon as evaluation words that interact with other terms. For instance, in the sentence *It was barely sufficient*, the evaluation of the combination is negative. Examples of nouns that act like shifters are *failure* and *neglect*. In the phrase ‘*failure to succeed*’, for example, the force of the meaning of *failure* transforms the positive valence of *succeed* into a negative property. The expression as a whole counts as negative.

The same observations can be made with respect to verbs like *fail*, *omit*, *neglect*.... They not only convey the information that something did not happen but also that the author was expecting it to happen and that this not borne out expectation has negative consequences as illustrated by *He stayed around* versus *He failed to leave*.⁶

3.1.4 Irony

Sometimes the contributions made by various lexical items combine in ways that cannot be accounted for in the ways described above. For example, in the ironic sentence *The very brilliant organizer failed to solve the problem*⁷, the extremely positive connotation of *very brilliant* is turned against itself by the meaning of the sentence. We account for this phenomena by assuming that in the lexicon *brilliant* will be marked as +2, *very* will increase the base valence of the expression to +3; *fail* will be marked as negative and the expression *solve the problem* will be marked positive. Evaluative terms under the scope of *fail*, such as *solve the problem* will be marked 0; entities whose existence is not denied by the use of *fail* but to whom failure is ascribed will turn negative. In this case, the base score was 0, however, *very brilliant* goes from positive to negative and *solve the problem* is neutralized, while *fail* remains negative. The adjusted score is: -4.

⁶ Often the use of *fail* leads to an indirect negative evaluation of the person to whom the failure is attributed. This can be exploited in irony (see below).

⁷ Note that when we add *even*, the situation changes again. The sentence is not necessarily ironic. Items under the scope of words like *even* are neutralized. So the sentence *Even the brilliant organizer failed to solve the problem* is scored -1 for *fail* only.

<i>brilliant</i>	+2		<i>Original valence is adjusted by</i>
			<i>very</i>
<i>very brilliant</i>	+3	-3	<i>adjusted because of fail</i>
<i>failed</i>	-1	-1	
<i>solve the problem</i>	+1	<u>0</u>	<i>neutralized by fail</i>
<i>total score:</i>		-4	

Table 2. Valence calculation for “The very brilliant organizer failed to solve the problem.”

3.2 Discourse Based Contextual Valence Shifters

3.2.1 Connectors

Connectors such as *although*, *however*, *but*, *on the contrary*, *notwithstanding* etc. can both introduce information, and act on information elsewhere in the text to mitigate the force of that information⁸. For example, take the sentence *Although Boris is brilliant at math, he is a horrible teacher*. While the statement *Boris is brilliant at math* positively assesses Boris’ math skills, the force of *although* combined with the negative assessment in the sentence’s main clause *he is a horrible teacher* effectively negates the positive force of the evaluation as applied to Boris. In computing the author’s attitude towards Boris, therefore, the effect of *although* is to neutralize the effect of the positive assessment, resulting in a negative assessment score for the sentence. Let’s follow that along step-by-step to make the claim clear:

<i>Although Boris is brilliant at math, he is a horrible teacher.</i>			
Base valence of terms:		Adjusted computation:	
<i>brilliant</i>	+2	<i>(Although) brilliant</i>	0
<i>horrible</i>	<u>-2</u>	<i>horrible</i>	<u>-2</u>
total score:	0	total score:	-2

Table 3. Example of valence adjustment based on discourse connective.

In this example we also see how the micro organization of the discourse makes a difference: the positive effect of *brilliant* is encapsulated in the embedded clause and does not contribute to the evaluation of the larger unit.

3.2.2 Discourse Structure and Attitude Assessment

A third discourse level valence adjuster included in this paper concerns *discourse structure* itself. There are two basic discourse relations of interest to us here: *lists* and *elaborations*. Some discourse constituents are linked to others in a list in which each constituent encodes a similar relationship to some more general concept and other constituents that give more detailed information of some sort about material encoded in constituents preceding them in the linear organization of the text. These earlier constituents structurally dominate the elaborating constituents (Grosz and Sidner, 1986; Mann and Thompson, 1988; Polanyi and Scha, 1984). Of

⁸ As was noticed by Hatzivassiloglou and McKeown (1997), the construction *Adj1 but Adj2* can be used to determine the valency of one adjective if the valency of the other one is known.

interest to us here is how base lexical valence scores are modified by their position in a hierarchical discourse structure.

In an Elaboration, a constituent gives more detail about another constituent which is in a structurally accessible position in a discourse stack. For instance in *John walks a lot. Last month he walked 25 miles on Tuesdays*, the second sentence illustrates the concept expressed in the dominating sentence. When valence information is introduced in a dominating sentence, the elaborations reinforce its effects. For example, lexical valence information is introduced by the use of *terrific* in the dominating sentence in the following passage:

*John is a **terrific*** athlete. Last week he walked 25 miles on Tuesdays. Wednesdays he walked another 25 miles. Every weekend he hikes at least 50 miles a day.*

Each of the dominated constituents is itself neutrally valenced. However, in this text, each is an example of John's terrific athleticism. Therefore, the positive valence of *terrific* is inherited by each subsequent new example. Effectively, the force for this one instance of the positively valenced term *terrific* as applied to *John* is greatly strengthened when the sentence is treated in its discourse context rather than as an independent expression.

3.2.3 Multi-entity Evaluation

Up to now we have looked at the effects that context can have on the evaluation of one single entity. But in most complex documents a wide variety of entities are discussed – some of which might be evaluated positively and others negatively. For example, a product reviewer discusses one negative aspect of a product extensively in a review which was otherwise very positive about many other features. In this case, it would be incorrect to assume that the reviewer was negative towards the product because of having described one negative feature in some detail. In such a case, simple methods of comparing the number of positive terms versus the number of negative terms could result in a faulty assessment of the reviewer's attitude towards the product. No simple correlation need obtain between the length at which a particular aspect of situation is discussed and the weight that discussion plays in an overall assessment.

3.2.4 Genre and Attitude Assessment

The assessment of author attitude may be complexly related to the genre of the communication in which valence marked terms occur. For example, any use of evaluative language in a document in which such assessments seldom occur will carry more weight than would otherwise be the case. Similarly, the presence of valence carrying items in a text by an author or found in a text type associated with the use of highly evaluative language may carry less weight. As we show below, assessing attitude in a document in which there are various participants "speaking" in a text can be an issue as well.

3.2.5 Reported Speech

Take the sentence *Mary was a slob*. The base valence of this sentence is -1, since *slob* is a negatively valenced term. Now, consider, *John said that Mary was a slob*. Here the author asserts that John said something unflattering about Mary, not that the author accepts John's assessment. However, information later in the text could force its inclusion as in *John said that Mary was a slob and he is right*. In this case, the negative valence attached to *slob* will be counted along with the positive valence of *right*. To illustrate consider this text:

*The utilities argue that they **performed well**⁺. But the public still remembers those **miserable**⁻ **rotten**⁻ nights.*

Both *argue* and *remembers* are Reported Speech and Thought operators. Therefore, the valence of the reported material is not ascribed to the author (Wiebe, Wilson and Bell, 2001; Wiebe, Wilson, Bruce, Bell and Martin, 2004) but to the utilities and the public respectively. The positive and the negative valences do not cancel each other out. The text is not neutral; it is positive in relation to utilities and and negative in relation to public. We need two different counters one for the utilities and one for the public.

Notice also that the weight of both valences is not equal for the larger unit composed of both sentences. By using *But* the author chooses to give more weight to the second point of view reported as this point in the text. A sensitive weighing scheme could be devised to reflect these complex facts (Riloff and Wiebe, 2003).

3.2.6 Subtopics

Sometimes it is possible to split a longer document into subtopics. The point of view of the author can then be made relative to each subtopic. Take for instance the following artificially constructed short text.

*Our yearly overview of the situation in Ubitopia.
The economic situation is more than **satisfactory**⁺. The leading indicators show a **rosy**⁺ picture.
The manufacturing sector is **booming**⁺. Exports have **exceeded**⁺ the **wildest expectations**⁺.
When one looks at the human rights picture, one is struck by the increase in **arbitrary**⁻ arrests, by needless **persecution**⁻ of **helpless**⁻ citizens and increase of police **brutality**⁻.*

In a text like this, one could link the positive and the negative attitudes to the two subtopics, the economy and the human rights situation. In most cases, this will not be as easy as is the case here, and, even if the text can be clearly divided into subtopics, it is not necessarily the case that all subtopics contribute equally to the overall impression that a text makes. One factor that will influence their contribution is genre, which we discuss next.

3.2.7 Genre Constraints

Movie reviews have been a focus of attention in the document classification community for some time. These texts are known to be notoriously difficult to work with using existing techniques (Pang, Lee and Vaithyanathan, 2002). Problems arise because these texts are composed of two types of information: information about the events and situations in the story and information about the film which has been created to tell the story. Since the question one is interested in primarily interested in having answered by a movie review is *Is this a good movie?* and since the review is prepared by the reviewer to answer this question, it is necessary to separate the description of the entities pertaining to the story from the description of the entities pertaining to the production. Only the valence scores of the entities pertaining to the production should be considered in ascertaining if the review is positive or negative.

Reviews of films loosely follow a set of genre conventions that can be mined for factors which can influence basic valence assignment. For example, movie reviews are often constructed as a quasi-interaction between author and reader. Comments in or about the first or second person reflect information about the film since neither reader nor author are characters in the film. Positional

information can also be important: comments at the very beginning or very end of a review are accorded more weight than remarks in less prominent positions.

Let's consider an artificially constructed example based on an excerpt from a movie review taken from MRDb website used in Pang, Lee and Vaithyanathan (2002)⁹:

*This film should be **brilliant**⁺. The characters are **appealing**⁺. ... Stallone plays a **happy**⁺, **wonderful**⁺ man. His **sweet**⁺ wife is **beautiful**⁺ and **adores**⁺ him. He has a **fascinating**⁺ **gift**⁺ for **living life fully**⁺. It sounds like a **great**⁺ plot, however, the film is a **failure**⁻*

	Adjusted Score
<i>This film should be brilliant⁺.</i>	0 <i>brilliant</i> within scope of <i>should</i> is 0
<i>The characters are appealing⁺.</i>	0 <i>appealing</i> elaboration under <i>should</i> 0
<i>Stallone plays a happy⁺, wonderful⁺ man. His sweet⁺ wife is beautiful⁺ and adores⁺ him. He has a fascinating⁺ gift⁺ for living life fully⁺.</i>	0 <i>Happy, wonderful, sweet</i> etc. all refer to storyworld entities and thus are not counted.
<i>It sounds like a great⁺ story,</i>	-1 <i>However</i> reverses + valence of <i>great</i>
<i>however, the film is a failure⁻</i>	-1 <i>failure</i> is -1
Total Score:	-2

Table 4. Valence calculation for the movie review.

The adjusted score is -2. The review is negative.

In some cases, we should be able to exploit genre constraints in determining the attitude of authors towards the entities created in the documents. But to do this computationally, the structure that genres impose on documents needs to be determined automatically. This is not yet possible.

4. Conclusion

We have shown that even when the author attention is restricted to one topic/entity/fact, lexical items in a discourse context will interact with one another. An author's attitude cannot be calculated based on individual items. We proposed a calculation of local interactions that improves upon the results of current approaches based on simple counts. We also argued that valence calculation is critically affected by discourse structure. In addition, we discussed cases in which a document describes more than one entity/topic/fact. We showed that, in these cases, the calculation of point of view must be done with respect to each entity separately and must take into account higher order factors such as genre that influence document structure

Taken together, these considerations argue strongly that calculating author attitude must be based on a finer grained analysis of the text on all levels than has been previously proposed.

⁹ The authors explain that their context insensitive evaluative lexical methods fail on texts in which the author sets up a deliberate contrast to an expected position. They cannot deal with the mismatch between the base valence of the term and the author's usage.

5. Bibliography

- Das, S. and Chen, M. (2001) *Yahoo! For Amazon: extracting market sentiment from stock message boards*. Paper presented in the 8th Asia Pacific Finance Association Annual Conference, Bangkok, Thailand.
- Edmonds, P. and Hirst, G. (2002) *Near-synonymy and lexical choice*. *Computational Linguistics*, 28 (2), 105–144.
- Grosz, B. and Sidner, C. (1986) *Attention, Intention, and the Structure of Discourse*. *Computational Linguistics*, 12 (3), 175–204.
- Hatzivassiloglou, V. and McKeown K. (1997) *Predicting the Semantic Orientation of Adjectives*. *Computational Linguistics*, 174–181.
- Hatzivassiloglou, V. and Wiebe, J. (2000) *Effects of Adjective Orientation and Gradability on Sentence Subjectivity*. Paper presented at Coling 2000. Saarbrücken.
- Mann, W. and Thompson, S. (1988) *Rhetorical Structure Theory: Toward a functional theory of text organization*. *Text* 8 (3), 243–281.
- Pang, B. Lee, L. and Vaithyanathan, V. (2002) *Thumbs up? Sentiment Classification using Machine Learning Techniques*. In *Proceedings of the 2002 Conference on Empirical Methods on Natural Language Processing*. 79–86.
- Polanyi, L. and Scha, R. (1984) *A Syntactic Approach to Discourse Semantics*. Paper presented at 6th International Conference on Computational Linguistics, Stanford, CA.
- Riloff, E. and Wiebe, J. (2003) *Learning Extraction Patterns for Subjective Expressions*. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*. 105–112.
- Tong, R. (2001) *An operational system for detecting and tracking opinions in on-line discussion*. Paper presented at the SIGIR Workshop on Operational Text Classification.
- Turney, P. and Littman, M. (2002) *Unsupervised Learning of Semantic Orientation from a Hundred-Billion-Word Corpus*. Technical Report, National Research Council Canada, Institute for Information Technology, ERB-1094, NRC-44929.
- Wiebe, J., Wilson, T., and Bell, M. (2001) *Identifying Collocations for Recognizing Opinions*. Paper presented at the ACL/EACL '01 Workshop on Collocation, Toulouse, France, July 2001.
- Wiebe, J., Wilson, T., Bruce, R., Bell, M., and Martin, M. (2004) *Learning subjective language*. *Computational Linguistics* 30 (3).