

# Learning Lexico-Functional Patterns for First-Person Affect

Lena Reed, Jiaqi Wu, Shereen Oraby  
Pranav Anand and Marilyn Walker

University of California Santa Cruz

{lreed, jwu64, soraby, panand, mawalker}@ucsc.edu

## Abstract

Informal first-person narratives are a unique resource for computational models of everyday events and people’s affective reactions to them. People blogging about their day tend not to explicitly say *I am happy*. Instead they describe situations from which other humans can readily infer their affective reactions. However current sentiment dictionaries are missing much of the information needed to make similar inferences. We build on recent work that models affect in terms of lexical predicate functions and affect on the predicate’s arguments. We present a method to learn proxies for these functions from first-person narratives. We construct a novel fine-grained test set, and show that the patterns we learn improve our ability to predict first-person affective reactions to everyday events, from a Stanford sentiment baseline of .67F to .75F.

## 1 Introduction

Across social media, thousands of posts daily take the form of informal FIRST-PERSON NARRATIVES. These narratives provide a rich resource for computational modeling of how people feel about the events they report on. Being able to reliably predict the affect a person may feel towards events they encounter has a range of potential applications, including monitoring mood and mental health (Isaacs et al., 2013) and getting conversational assistants to respond appropriately (Bowden et al., 2017). Moreover, as these narratives are told from the perspective of a protagonist, this research could be used to understand other types of protagonist-framed narratives, like those in fiction.

We are interested in the opinions that a protagonist has, not the author per se. This is sometimes referred to as internal sentiment or self reflective sentiment. While in many situations that is overlaid with the author’s opinions, in first-personal narratives, because the author is the protagonist, the two perspectives align. Here, we use the term *affect* to reference this protagonist-centered notion of opinion.

A central obstacle to reliable affect prediction is that that people tend not to *explicitly* flag their affective state, by saying *I am happy*. Large-scale sentiment dictionaries focus on compiling lexical items that bear a consistent affect all on their own (Wilson et al., 2005). But people tend to describe situations, such as *My friend bought me flowers*, or *I got a parking ticket*, from which other humans can readily infer their *implicit* affective reactions.

One approach to this problem aims to directly learn units larger than a lexical item that reliably bear some marker of polarity or emotion (Vu et al., 2014; Li et al., 2014; Ding and Riloff, 2016; Goyal et al., 2010; Russo et al., 2015; Kiritchenko et al., 2014; Reckman et al., 2013).

$X$	$Y$	$E_{hve}$	$E_{tck}$	Example
+	+	+	-	<i>I have a new kitten.</i>
+	-	-	+	<i>I got a parking ticket.</i>
-	+	-	+	<i>My rival got a prize.</i>
-	-	+	-	<i>My rival got a reprimand.</i>

$X$  have/lack  $Y$

Table 1: Functions for verbs of possession.

Another approach aims to model the speaker’s affect to an event compositionally, e.g. Anand and Reschke (2010) (A&R) proposed that the affect a lexical predicate communicates should be modeled as an n-ary function, taking as inputs the affect that the speaker bears towards each partici-

pant. Table 1 contains A&R’s functions for verbs of possession: a state in which  $X$  *has*  $Y$  or  $X$  *lacks*  $Y$  does not convey a clear affect unless we know what the speaker thinks of both  $X$  and  $Y$ . If the speaker has positive affect toward both  $X$  and  $Y$  (Row 1), then we infer that her attitude toward the event is positive, but if either is negative, then we infer that the speaker is negative toward the event. Similarly, Rashkin et al. (2016) represent the typical affect communicated by particular predicates via connotation frames. Here we are finding the internal sentiment of the speaker, or, as Rashkin et al. refer to it, the “mental state” of the speaker.

Inspired by A&R’s framework, our work learns lexico-functional patterns (patterns involving lexical items or pairs of lexical items in specific grammatical relations that we show to capture functor-argument relations in A&R’s sense), about the effects of combining particular arguments with particular verbs (event types) from first-person narratives. Our novel observation is that learning these compositional functions is greatly simplified in the case of first-person affect. People bear positive affect to themselves, so sentences with first-person elements, e.g. *I/we/me*, reduce the problem for an approach like A&R’s to learning the polarity that results from composing the verb with *only one* of its arguments, i.e. only Rows 1, 2 in Table 1 need to be learned for first person subjects. First-person narratives are full of such sentences. See Table 2. We show that the learned patterns are often consonant with A&R’s predictions, but are richer, including e.g. many private state descriptions (Wiebe et al., 2004; Wiebe, 1990).

Positive Sentences
We had a marvelous visit and drank coffee and ate homemade chocolate chip cookies.
Now, I could swim both froggy and free style swimming!!
Negative Sentences
But last week, he said that he doesn’t know if he has the same feelings for me anymore.
I didn’t want to lose him.

Table 2: Sentences from the training data

In addition, we demonstrate that these lexico-functional patterns improve the performance of several off-the-shelf sentiment analyzers. We show that Stanford sentiment (Socher et al., 2013) has a best performance of 0.67 macro F on our test set. We then supplement it with our learned patterns and demonstrate significant improvements. Our final ensemble achieves 0.75 F on the test set.

We discuss related work in more detail in Sec. 5.

## 2 Bootstrapping a First-Person Sentiment Corpus

We start with a set of first-person narratives (we-blogs) drawn from the Spinn3r corpus, that cover a wide range of topics (Burton et al., 2009; Gordon and Swanson, 2009). To reduce noise, we restrict the blogs to those from well-known blogging sites (Ding and Riloff, 2016), and select 15,466 stories whose length ranges from 225 to 375 words.

Pattern Template	Class	Example Instantiations
<subj> ActVP	neg	<I> cry at the thought of it and I’m crying now.
<subj> ActInfVP	pos	<I> got to swim from the boat to a little sandbar.
<subj> AuxVP Dobj	neg	Yesterday <it> was my 1st molar, today it’s my 2nd molar.
ActVP <dobj>	pos	As I bake often, I have used <several different kinds of recipes>.
PassInfVP <dobj>	pos	When we arrived at the Embarcadero, we were surprised to find <a music festival taking place>...
Subj AuxVP <dobj>	neg	Our relationship was <non existent> for over a year after that.
NP Prep <np>	neg	he hurt me countless times but I still forgave him and i still tried to prove to him that I did care for <him>.
ActVP Prep <np>	neg	I didn’t think anything of it until I thought about when he cheated on <me>.
InfVP Prep <np>	pos	...my friend from college who was so generous to offer his place to <us>...

Table 3: AutoSlog-TS Templates and Example Instantiations

We hand-annotate a set of 477 positive and 440 negative stories, and use these to bootstrap a larger set of 1,420 negative and 2,288 positive stories. To bootstrap, we apply AutoSlog-TS, a weakly supervised pattern learner that only requires training sets of stories labeled broadly as POSITIVE or NEGATIVE (Riloff, 1996; Riloff and Wiebe, 2003). AutoSlog uses a set of syntactic templates to define different types of linguistic expressions. The left-hand side of Table 3 lists examples of AutoSlog patterns and the right-hand side illustrates a specific lexical-syntactic pattern that corresponds to each general pattern template, as instantiated in first-person stories.<sup>1</sup> When bootstrapping a larger positive and negative story corpus, we use the whole story, not just the first person sentences.

The left-hand-side of Table 3 shows that the learned patterns can involve syntactic arguments

<sup>1</sup>The examples are shown as general expressions for readability, but the actual patterns must match the syntactic constraints associated with the pattern template.

of the verbal predicate, which means that these patterns are proxies for one column of verbal function tables like those in Table 1. However, they can also include verb-particle constructions, such as *cheated on*, or verb-head-of-preposition constructions. In each case though, because these patterns are localized to a verb and only one element, they allow us to learn highly specific patterns that could be incorporated into a dictionary such as +-Effect (Choi and Wiebe, 2014). AutoSlog simultaneously harvests both (syntactically constrained) MWE patterns and more compositionally regular verb-argument groups at the same time.

AutoSlog-TS computes statistics on the strength of association of each pattern with each class, i.e.  $P(\text{POSITIVE} | p)$  and  $P(\text{NEGATIVE} | p)$ , along with the pattern’s overall frequency. We define three parameters for each class:  $\theta_f$ , the frequency with which a pattern occurs,  $\theta_p$ , the probability with which a pattern is associated with the given class and  $\theta_n$ , the number of patterns that must occur in the text for it to be labeled. These parameters are tuned on the dev set (Riloff, 1996; Oraby et al., 2015; Riloff and Wiebe, 2003).

To bootstrap a larger corpus, we want settings that have lower recall but very high precision. We select  $\theta_p = 0.7$ ,  $\theta_f = 10$  and  $\theta_n = 3$  for the positive class and  $\theta_p = 0.85$ ,  $\theta_f = 10$  and  $\theta_n = 4$  for the negative class for bootstrapping.

### 3 Experimental Setup

Our experimental setup involves first creating a corpus of training and test **sentences**, then applying AutoSlog-TS a second time to learn linguistic patterns. We then set up methods for cascading classifiers to explore whether ensemble classifiers improve our results.

**Training Set:** From the bootstrapped set of stories, we create a corpus of sentences. A critical simplifying assumption of our method is that a multi-sentence story can be labelled as a whole as positive or negative, and that each of its sentences **inherit** this polarity. This means we can learn the polarity of events in such narratives from their (noisy) inherited polarity without labelling individual sentences. Our training set consists of 46,255 positive and 25,069 negative sentences.

**Test Set:** We create the test set by selecting 4k random first-person sentences. First-person sentences either contain an explicit first person marker, i.e. *we* and *my* or start with either a progressive verb or pleonastic *it*. To collect gold labels, we designed a qualifier and a HIT for Mechanical Turk, and put

these out for annotation by 5 Turkers, who label each instance as positive, negative, or neutral. To ensure the high quality of the test set, we select sentences that were labelled consistently positive or negative by 4 or 5 Turkers. We collected 1,266 positive and 1,440 negative sentences.

**Dev Set:** We created the dev set using the same method as the test set, having Turkers annotate 2k random first-person sentences. We collected 498 positive and 754 negative sentences. The 4k test and dev sentences available for download at <https://nlds.soe.ucsc.edu/first-person-sentiment>.

**AutoSlog First-Person Sentence Classifier.** In order to learn new affect functions, we develop a second **sentence-level** classifier using AutoSlog-TS. We run AutoSlog over the training corpus, using the dev set to tune the parameters  $\theta_f$ ,  $\theta_p$  and  $\theta_n$  (?), in order to maximize macro F-score. Our best parameters on the dev set for positive is  $\theta_f=18$ ,  $\theta_p=0.85$  and  $\theta_n=1$  and for negative is  $\theta_f=1$ ,  $\theta_p=0.5$  and  $\theta_n=1$ . We specify that if the sentence is in both classes we rename it as neutral. We will refer to this classifier as the AutoSlog classifier.

**Baseline First-Person Sentence Classifiers.** Our goal is to see whether the knowledge we learn using AutoSlog-TS complements existing sentiment classifiers. We thus experiment with a number of baseline classifiers: the default SVM classifier from Weka with unigram features (Hall et al., 2005), a version of the NRC-Canada sentiment classifier (Mohammad et al., 2013), provided to us by Qadir and Riloff (2014), and the Stanford Sentiment classifier (Socher et al., 2013).

**Retrained Stanford.** The Stanford Sentiment classifier is a based on Recursive Neural Networks, and trained on a compositional Sentiment Treebank, which includes fine-grained sentiment labels for 215,154 phrases from 11,855 sentences from movie reviews. It can accurately predict some compositional semantic effects and handle negation. However since it was trained on movie reviews, it is likely to be missing labelled data for some common phrases in our blogs. Thus we also retrained it (RETRAINED STANFORD) on high precision phrases from AutoSlog extracted from our training data of positive and negative blogs. This provides 67,710 additional phrases, including 58,972 positive phrases and 8,738 negative phrases. The retrained model includes **both** the labels from the original Sentiment Treebank and the AutoSlog high precision phrases.

## 4 Results and Analysis

We present our experimental results and analyze the results in terms of the lexico-functional linguistic patterns we learn.

**Baseline Classifiers.** Rows 1-3 of Table 4 show the results for the three baselines, in terms of F-score for each class and the macro F. Stanford outperforms both NRC and SVM, but misses many cases of positive sentiment.

**AutoSlog Classifier.** Row 4 of Table 4 shows the results for the AutoSlog classifier. Although AutoSlog itself does not perform highly, the patterns that it learns represent a different type of knowledge than what is contained in many sentiment analysis tools. We therefore hypothesized that a cascading classifier, which supplements one of the baseline sentiment classifiers with the lexico-functional patterns that AutoSlog learns might yield higher performance.

**Retrained Stanford.** Row 5 of Table 4 shows the results for RETRAINED STANFORD. The F-scores for RETRAINED STANFORD are almost identical to the standard Stanford classifier. This may be because our data is a small percentage of the entire number of phrases used in training Stanford. Although RETRAINED STANFORD prioritizes our phrases, it would not make sense to remove the original training data.

**Cascading Classifiers.** We implement cascading classifiers to test our hypothesis. The cascade classifier has primary and secondary classifiers, and we invoke the secondary classifiers only if the primary assigns a prediction of *neutral* to a test instance, which reflects the lack of sentiment-bearing lexical items. We also have a cascade classifier with a tertiary classifier, which is invoked in the same fashion as the secondary classifier after the primary and secondary classifiers have been run. The cascading classifiers are named in the order the classifier is employed, primary, secondary or primary, secondary, tertiary. For our cascading classifiers, we combine our baseline classifiers (NRC and Stanford), with our AutoSlog classifier. We do not use SVM as a primary classifier since it has no neutral label. The results for the cascading experiments are in Rows 6-9 of Table 4.

Cascading NRC and AutoSlog provides the best performance, improving both the positive and negative classes, for a macro F of 0.71. This shows that the learned implicit polarity information from AutoSlog improves NRC’s performance.

Since our best two-classifier cascade comes

from combining NRC and AutoSlog, we also test a cascade that adds Stanford or SVM. We achieve our best macro F of 0.75 for the combination with SVM.

	Classifier	Pos F1	Neg F1	Macro F
1	SVM	0.66	0.60	0.64
2	NRC	0.58	0.69	0.64
3	Stanford	0.54	0.73	0.67
4	AutoSlog (ASlog)	0.11	0.68	0.53
5	Retrained Stanford	0.53	0.73	0.67
6	NRC, ASlog	0.60	0.78	0.71
7	Stanford, ASlog	0.55	0.76	0.70
8	NRC, ASlog, Stanford	0.64	<b>0.79</b>	0.74
9	NRC, ASlog, SVM	<b>0.70</b>	0.78	<b>0.75</b>

Table 4: Test Set Results

**Analysis and Discussion.** Here we discuss how the patterns we learned from AutoSlog can supplement the knowledge encoded in current sentiment classifiers, and in newly evolving sentiment resources (Goyal et al., 2010; Choi and Wiebe, 2014; Balahur et al., 2012; Ruppenhofer and Brandes, 2015).

POS PATTERNS	Basic Entailment
HAVE_FUN	property
HAVE_PARTY	possession
HEADED_FOR	location
NEG PATTERNS	Basic Entailment
HAVE_CANCER	property
LOST	possession
NOT_COME_HOME	location
NOT_GOING_KILL	existence

Table 5: Highly predictable AutoSlog extracted case frames and functional description

Tables 5 and 6 illustrate several learned lexico-functional patterns for positive events used in the AutoSlog classifier. The patterns shown in Table 5 are predicted by A&R’s framework, some functions of which can be seen in Table 1. For example, we find a range of basic state descriptions (*have.party*, *have.cancer*) whose basic entailment category is either *possessive* or *property* state. Since  $E_{have}$  is positive for a first-person subject only if the object is positive, and negative if the object is negative, we predict that *parties* are good to possess and that *cancer* is a bad property to have. In this way, we can recruit the existing function for *have* to induce new positive or negative things to “possess.” In line with A&R’s claims, many events are identified with their final results: *headed for* results in being at a desired location, while *not coming home* results in something failing to be at a desired location. We find it a welcome result that our semi-supervised methods

yield patterns that correspond to the A&R classes, thus validating our suspicion that first-person sentences furnish a simplifying test ground for discovering functional patterns in the wild.

However, many patterns are not covered by A&R’s general classes, see Table 6. Looking first at verbs, one major correlation is between positive classes and public events and negative classes and private states. Verbs extracted from the positive class tend to be eventive and agentive describing more dynamic activities and interactions, such as *played*, *swim*, *enjoyed*, and *danced*. Even many positive *have* uses are light verbs describing an activity such as *have lunch*.

Description	POS PATTERNS
activities	HAVE_DINNER, HAVE_WEDDING
success	GOT_SEE, WENT_WELL
planning	HEADED_FOR, SET_UP
free time	HAVE_TIME, TIME_WITH
social bonding	PICTURE_OF, OLD_FRIENDS
Description	NEG PATTERNS
activities	HAVE_X-RAY, GET_EXAM
knowledge	REALIZE, NOT_KNOW_WHAT
unmet desire	WANTS, NEED_MONEY
social bonding	NOT_TRUST

Table 6: Highly predictable AutoSlog case frames outside A&R’s functional system

Verbs from the negative class are strikingly different. They are very often stative, where the author is the experiencer (cognitive subject) of that private state. While this state vs. event distinction is not one existing computational models of sentiment or affect discuss explicitly, it replicates a finding that consistently emerges in clinical psychology, one that is explicitly argued for in cognitive-behavioral accounts of the mood that particular activities evoke (Lewinsohn et al., 1985; MacPhillamy and Lewinsohn, 1982; Russo et al., 2015). In addition, Table 6 reveals several novel result state categories. The success, planning, and unmet desire frames are all ultimately about goal-fulfillment (or lack thereof). While the success and unmet desire cases could be understood as having or lacking something, the planning cases indicate steps achieved toward a desired end-state. Previous work on learning affect from eventuality descriptions has largely focused on actions. Our results indicate that private state descriptions are another rich source of evidence.

## 5 Related Work

Previous work learns phrasal markers of implicit polarity via bootstrapping from large-scale text sources, e.g. Vu et al. (2014)

learn emotion-specific event types by extracting *emotion, event* pairs on Twitter. Li et al. (2014) uses Twitter to bootstrap ‘major life events’ and typical replies to those events.

Ding and Riloff (2016) extract subj-verb-obj triples from blog posts. They then apply label propagation to spread polarity from sentences to events. However, the triples they learn do not focus on first-person experiencers. They also filter private states out of the verbs used to learn their triples, whereas we have found that verbs relating to private states such as *need*, *want* and *realize* are important indicators of first-person affect.

Balahur et al. (2012) use the narratives produced by the ISEAR questionnaire (Scherer et al., 1986) for first-person examples of particular emotions and extract sequences of subject-verb-object triples, which they annotate for basic emotions.

Recent work has built on this idea, and developed methods to automatically expand Anand & Reschke’s verb classes to create completely new lexical resources (Balahur et al., 2012; Choi and Wiebe, 2014; Deng et al., 2013; Deng and Wiebe, 2014; Ruppenhofer and Brandes, 2015). Choi & Wiebe’s work comes closest to ours in trying to induce (not annotate) lexical functions, but we attempt to infer these from stories directly, whereas they use a structured lexical resource.

## 6 Conclusion

We show that we can learn lexico-functional linguistic patterns that reliably predict first-person affect. We constructed a dataset of positive and negative first-person experiencer sentences and used them to learn such patterns. We then showed that the performance of current sentiment classifiers can be enhanced by augmenting them with these patterns. By adding our AutoSlog classifier’s results to existing classifiers we were able to improve from a baseline 0.67 to 0.75 Macro F with a cascading classifier of NRC, AutoSlog and SVM. In addition, we analyze the linguistic functions that indicate positivity and negativity for the first person experiencer, and show that they are very different. In first-person descriptions, positivity is often signaled by active participations in events, while negativity involves private states. In future work, we plan to explore the integration of these observations into sentiment resources such as the +-Effect lexicon (Choi and Wiebe, 2014). We plan to apply these high precision first-person lexical patterns beyond blog data and with other person-marking.

## References

- Pranav Anand and Kevin Reschke. 2010. Verb classes as evaluativity functor classes. *Proceedings of Verb* pages 98–103.
- Alexandra Balahur, Jesus M. Hermida, and Andres Montoyo. 2012. Building and exploiting emotinet, a knowledge base for emotion detection based on the appraisal theory model. *IEEE Trans. Affect. Comput.* 3(1):88–101.
- Kevin Bowden, Shereen Oraby, Amita Misra, Jiaqu Wu, Stephanie Lukin, and Marilyn Walker. 2017. Data-driven dialogue systems for social agents. In *International Workshop on Spoken Dialogue Systems*.
- Kevin Burton, Akshay Java, and Ian Soboroff. 2009. The icwsm 2009 spinn3r dataset. In *Proceedings of the Annual Conference on Weblogs and Social Media (ICWSM)*.
- Yoonjung Choi and Janyce Wiebe. 2014. +/-effectwordnet: Sense-level lexicon acquisition for opinion inference. In *EMNLP*.
- Lingjia Deng, Yoonjung Choi, and Janyce Wiebe. 2013. Benefactive/malefactive event and writer attitude annotation. In *ACL*. pages 120–125.
- Lingjia Deng and Janyce Wiebe. 2014. Sentiment propagation via implicature constraints. In *EACL*. pages 377–385.
- Haibo Ding and Ellen Riloff. 2016. Acquiring knowledge of affective events from blogs using label propagation. In *AAAI*.
- Andrew Gordon and Reid Swanson. 2009. Identifying personal stories in millions of weblog entries. In *Third International Conference on Weblogs and Social Media, Data Challenge Workshop, San Jose, CA*.
- Amit Goyal, Ellen Riloff, and Hal Daumé III. 2010. Automatically producing plot unit representations for narrative text. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. pages 77–86.
- M. Hall, F. Eibe, G. Holms, B. Pfahringer, P. Reutemann, and I. Witten. 2005. The weka data mining software: An update. *SIGKDD Explorations* 11(1).
- Ellen Isaacs, Artie Konrad, Alan Walendowski, Thomas Lennig, Victoria Hollis, and Steve Whittaker. 2013. Echoes from the past: how technology mediated reflection improves well-being. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, pages 1071–1080.
- Svetlana Kiritchenko, Xiaodan Zhu, and Saif M. Mohammad. 2014. Sentiment analysis of short informal texts. *J. Artif. Int. Res.* 50(1):723–762.
- Peter M Lewinsohn, Robin M Mermelstein, Carolyn Alexander, and Douglas J MacPhillamy. 1985. The unpleasant events schedule: A scale for the measurement of aversive events. *Journal of Clinical Psychology* 41(4):483–498.
- Jiwei Li, Xun Wang, and Eduard Hovy. 2014. What a nasty day: Exploring mood-weather relationship from twitter. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, New York, NY, USA, CIKM '14, pages 1309–1318. <https://doi.org/10.1145/2661829.2662090>.
- Douglas J MacPhillamy and Peter M Lewinsohn. 1982. The pleasant events schedule: Studies on reliability, validity, and scale intercorrelation. *Journal of Consulting and Clinical Psychology* 50(3):363.
- Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. Nrc-canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of the seventh international workshop on Semantic Evaluation Exercises (SemEval-2013)*. Atlanta, Georgia, USA.
- Shereen Oraby, Lena Reed, Ryan Compton, Ellen Riloff, Marilyn Walker, and Steve Whittaker. 2015. And that's a fact: Distinguishing factual and emotional argumentation in online dialogue. In *Proceedings of the 2nd Workshop on Argumentation Mining*. pages 116–126.
- Ashequl Qadir and Ellen Riloff. 2014. Learning emotion indicators from tweets: Hashtags, hashtag patterns, and phrases. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics*. pages 1203–1209.
- Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016. Connotation frames: A data-driven investigation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 311–321. <http://www.aclweb.org/anthology/P16-1030>.
- Hilke Reckman, Cheyanne Baird, Jean Crawford, Richard Crowell, Linnea Micciulla, Saratendu Sethi, and Fruzsina Veress. 2013. teragram: Rule-based detection of sentiment phrases using sas sentiment analysis. In *SemEval@NAACL-HLT*.
- E. Riloff and J. Wiebe. 2003. Learning extraction patterns for subjective expressions. In *Proceedings of the 2003 conference on Empirical methods in natural language processing-Volume 10*. Association for Computational Linguistics, pages 105–112.
- Ellen Riloff. 1996. Automatically generating extraction patterns from untagged text. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2*. pages 1044–1049.

- Josef Ruppenhofer and Jasper Brandes. 2015. Extending effect annotation with lexical decomposition. In *6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis WASS 2015*. page 67.
- Irene Russo, Tommaso Caselli, and Carlo Strapparava. 2015. Semeval-2015 task 9: Clipeval implicit polarity of events. In *Proceedings of the 9th International Workshop on Semantic Evaluation, SemEval*.
- Klaus R. Scherer, Harald G. Wallbott, and Angela B. Summerfield. 1986. *Experiencing emotion : a cross-cultural study*. Cambridge University Press.
- Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. *Conference on Empirical Methods in Natural Language Processing (EMNLP) 2013*.
- Hoa Trong Vu, Graham Neubig, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2014. Acquiring a dictionary of emotion-provoking events. *EACL 2014* page 128.
- J. Wiebe, T. Wilson, R. Bruce, M. Bell, and M. Martin. 2004. Learning subjective language. *Computational linguistics* 30(3):277–308.
- Janyce M Wiebe. 1990. Identifying subjective characters in narrative. In *Proceedings of the 13th conference on Computational linguistics-Volume 2*. Association for Computational Linguistics, pages 401–406.
- T. Wilson, P. Hoffmann, S. Somasundaran, J. Kessler, J. Wiebe, Y. Choi, C. Cardie, E. Riloff, and S. Patwardhan. 2005. Opinionfinder: A system for subjectivity analysis. In *Proceedings of HLT/EMNLP on Interactive Demonstrations*. Association for Computational Linguistics, pages 34–35.