

EMOBANK: Studying the Impact of Annotation Perspective and Representation Format on Dimensional Emotion Analysis

Sven Buechel and Udo Hahn

Jena University Language & Information Engineering (JULIE) Lab
Friedrich-Schiller-Universität Jena, Jena, Germany
{sven.buechel, udo.hahn}@uni-jena.de
<http://www.julielab.de>

Abstract

We describe EMOBANK, a corpus of 10k English sentences balancing multiple genres, which we annotated with dimensional emotion metadata in the Valence-Arousal-Dominance (VAD) representation format. EMOBANK excels with a bi-perspectival and bi-representational design. On the one hand, we distinguish between writer’s and reader’s emotions, on the other hand, a subset of the corpus complements dimensional VAD annotations with categorical ones based on Basic Emotions. We find evidence for the supremacy of the reader’s perspective in terms of IAA and rating intensity, and achieve close-to-human performance when mapping between dimensional and categorical formats.

1 Introduction

In the past years, the analysis of affective language has become one of the most productive and vivid areas in computational linguistics. In the early days, the prediction of the semantic polarity (positiveness or negativeness) was in the center of interest, but in the meantime, research activities shifted towards a more fine-grained modeling of sentiment. This includes the extension from only two to multiple polarity classes or even real-valued scores (Strapparava and Mihalcea, 2007), the aggregation of multiple aspects of an opinion item into a composite opinion statement for the whole item (Schouten and Frasincar, 2016), and sentiment compositionality (Socher et al., 2013).

Yet, two important features of fine-grained modeling still lack appropriate resources, namely shifting towards psychologically more adequate models of emotion (Strapparava, 2016) and distinguishing between writer’s vs. reader’s perspec-

tive on emotion ascription (Calvo and Mac Kim, 2013). We close both gaps with EMOBANK, the first large-scale text corpus which builds on the Valence-Arousal-Dominance model of emotion, an approach that has only recently gained increasing popularity within sentiment analysis. EMOBANK not only excels with a genre-balanced selection of sentences, but is based on a *bi-perspectival* annotation strategy (distinguishing the emotions of writers and readers), and includes a *bi-representationally* annotated subset (which has previously been annotated with Ekman’s Basic Emotions) so that mappings between both representation formats can be performed. EMOBANK is freely available for academic purposes.¹

2 Related Work

Models of emotion are commonly subdivided into *categorical* and *dimensional* ones, both in psychology and natural language processing (NLP). Dimensional models consider affective states to be best described relative to a small number of independent emotional dimensions (often two or three): *Valence* (corresponding to the concept of polarity), *Arousal* (degree of calmness or excitement), and *Dominance*² (perceived degree of control over a situation); the VAD model. Formally, the VAD dimensions span a three-dimensional real-valued vector space as illustrated in Figure 1. Alternatively, categorical models, such as the six *Basic Emotions* by Ekman (1992) or the *Wheel of Emotion* by Plutchik (1980), conceptualize emotions as discrete states.³

In contrast to categorical models which were used early on in NLP (Ovesdotter Alm et al., 2005; Strapparava and Mihalcea, 2007), dimensional

¹<https://github.com/JULIELab/EmoBank>

²This dimension is sometimes omitted (the VA model).

³Both dimensional and categorical formats allow for numerical scores regarding their dimensions/categories.




Figure 1: The affective space spanned by the three VAD dimensions. As an example, we here include the positions of Ekman’s six Basic Emotions as determined by Russell and Mehrabian (1977).

models have only recently received increased attention in tasks such as word and document emotion prediction (see, e.g., Yu et al. (2015), Köper and Schulte im Walde (2016), Wang et al. (2016), Buechel and Hahn (2016)).

In spite of this shift in modeling focus, VA(D)-annotated corpora are surprisingly rare in number and small in size, and also tend to be restricted in reliability. ANET, for instance, comprises only 120 sentences designed for psychological research (Bradley and Lang, 2007), while Preoțiuc-Pietro et al. (2016) created a corpus of 2,895 English Facebook posts relying on only two annotators. Yu et al. (2016) recently presented a corpus of 2,009 Chinese sentences from various online texts.

As far as categorical models for emotion analysis are concerned, many studies use incompatible subsets of category systems, which limits their comparability (Buechel and Hahn, 2016; Calvo and Mac Kim, 2013). This also reflects the situation in psychology where there is still no consensus on a set of fundamental emotions (Sander and Scherer, 2009). Here, the VAD model has a major advantage: Since the dimensions are designed as being independent, results remain comparable dimension-wise even in the absence of others (e.g., Dominance). Furthermore, dimensional models are the predominant format for lexical affective resources in behavioral psychology as evident from the huge number of datasets available for a wide range of languages (see, e.g., Warriner et al. (2013), Stadthagen-Gonzalez et al. (2016), Moors et al. (2013) and Schmidtke et al. (2014)).

For the acquisition of VAD values from participant’s self-perception, the Self-Assessment Manikin (SAM; Lang (1980), Bradley and Lang (1994)) has turned out as the most important and

(to our knowledge) only standardized instrument (Sander and Scherer, 2009). SAM iconically displays differences in Valence, Arousal and Dominance by a set of anthropomorphic cartoons on a multi-point scale (see Figure 2).

While it is common for more basic sentiment analysis systems in NLP to map the many different possible interpretations of a sentence’s affective meaning into a single assessment (“its sentiment”), there is an increasing interest in a more fine-grained approach where emotion expressed by writers is modeled separately from emotion evoked in readers. An utterance like “Italy defeats France in the World Cup Final” may be completely neutral from the *writer’s* viewpoint (presumably a professional journalist), but is likely to evoke rather adverse emotions in Italian and French *readers* (Katz et al., 2007).

In this line of work, Tang and Chen (2012) examine the relation between the sentiment of microblog posts and the sentiment of their comments (as a proxy for reader emotion). Liu et al. (2013) model the emotion of a news reader jointly with the emotion of a comment writer using a co-training approach. This contribution was followed up by Li et al. (2016) who propose a two-view label propagation approach instead. However, to our knowledge, only Mohammad and Turney (2013) investigated the effects of these perspectives on annotation quality, finding differences in inter-annotator agreement (IAA) relative to the exact phrasing of the annotation task.

In a similar vein to the writer-reader distinction, identifying the *holder* or *source* of an opinion or sentiment also aims at describing the affective information entailed in a sentence in more detail (Wiebe et al., 2005; Seki et al., 2009). Thus, opinion statements that can directly be attributed to the writer can be distinguished from references to other’s opinions. A related task, the detection of *stance*, focuses on inferring the writer’s (dis)approval towards a given issue from a piece of text (Sobhani et al., 2016).

3 Corpus Design and Creation

The following criteria guided the data selection process of the EMOBANK corpus: First, complementing existing resources which focus on social media and/or review-style language (Yu et al., 2016; Quan and Ren, 2009), we decided to address several genres and domains of general English.

Corpus	Domain	Raw	Filtered
SE07	news headlines	1,250	1,192
MASC	blogs	1,378	1,336
	essays	1,196	1,135
	fiction	2,893	2,753
	letters	1,479	1,413
	newspapers	1,381	1,314
	travel guides	971	919
Sum		10,548	10,062

Table 1: Genre distribution of the raw and filtered EMOBANK corpus.

Second, we conducted a pilot study on two samples (one consisting of movie reviews, the other pulled from a genre-balanced corpus) to compare the IAA resulting from different annotation perspectives (e.g., the writer’s and the reader’s perspective) in different domains (see Buechel and Hahn (2017) for details). Since we found differences in IAA but the results remained inconclusive, we decided to annotate the whole corpus *bi-perspectively*, i.e., each sentence was rated according to both the (perceived) writer *and* reader emotion (henceforth, WRITER and READER).

Third, since many problems of comparing emotion analysis studies result from the diversity of emotion representation schemes (see Section 2), the ability to accurately map between such alternatives would greatly improve comparability across systems and boost the reusability of resources. Therefore, at least parts of our corpus should be annotated *bi-representationally* as well, complementing dimensional VAD ratings with annotations according to a categorical emotion model.

Following these criteria, we composed our corpus out of several categories of the *Manually Annotated Sub-Corpus of the American National Corpus* (MASC; Ide et al. (2008), Ide et al. (2010)) and the corpus of SemEval-2007 Task 14 *Affective Text* (SE07; Strapparava and Mihalcea (2007)). MASC is already annotated on various linguistic levels. Hence, our work will allow for research at the intersection of emotion and other language phenomena. SE07, on the other hand, bears annotations according to Ekman’s six Basic Emotion (see Section 2) on a [0, 100] scale, respectively. This collection of raw data comprises 10,548 sentences (see Table 1).

Given this large volume of data, we opted for a crowdsourcing approach to annotation. We chose CROWDFLOWER (CF) over AMAZON MECHANICAL TURK (AMT) for its quality control mechanisms and accessibility (customers of AMT,




Figure 2: The modified 5-point Self-Assessment Manikin (SAM) scales for Valence, Arousal and Dominance (row-wise). Copyright of the original SAM by Peter J. Lang 1994.

but not CF, must be US-based). CF’s main quality control mechanism rests on *gold questions*, items for which the acceptable ratings have been previously determined by the customer. These questions are inserted into a task to restrict the workers to those performing trustworthily. We chose these gold items by automatically extracting highly emotional sentences from our raw data according to JEMAS⁴, a lexicon-based tool for VAD prediction (Buechel and Hahn, 2016). The acceptable ratings were determined based on manual annotations by three students trained in linguistics. The process was individually performed for WRITER and READER with different annotators.

For each of the two perspectives, we launched an independent task on CF. The instructions were based on those by Bradley and Lang (1999) to whom most of the VAD resources developed in psychology refer (see Section 2). We changed the 9-point SAM scales to 5-point scales (see Figure 2) in order to reduce the cognitive load during decision making for crowdworkers. For the writer’s perspective, we presented a number of linguistic clues supporting the annotators in their rating decisions, while, for the reader’s perspective, we asked what emotion would be evoked in an *average* reader (rather than asking for the rater’s personal feelings). Both adjustments were made to establish more objective criteria for the exclusion of untrustworthy workers. We provide the instructions along with our dataset.

For each sentence, five annotators generated VAD ratings. Thus, a total of 30 ratings were gathered per sentence (five ratings for each of the three VAD dimensions and two annotation perspectives, WRITER and READER). Ten sentences were presented at a time. The task was available for work-

⁴<https://github.com/JULIELab/JEmAS>

ers located in the UK, the US, Ireland, Canada, Australia or New Zealand. The total annotation costs amounted to \$1,578.

Upon inspection of the individual judgments, we found that the VAD rating (1, 1, 1) was heavily overrepresented. We interpret this skewed coding distribution as a bias mainly due fraudulent responses since, from a psychological view, this rating is highly improbable (Warriner et al., 2013). Accordingly, we decided to remove all of these ratings (about 10% for each of the tasks; the ‘Filtered’ condition in Table 1) because these annotations would have inserted a systematic bias into our data which we consider more harmful than erroneously removing a few honest outliers. For each sentence with two or more remaining judgments, its final emotion annotation is determined by averaging these valid ratings leading to a total of 10,062 sentences bearing VAD values for *both* perspectives (see Table 1).

This makes EMOBANK to the best of our knowledge by far the largest corpus for dimensional emotion models and, with the exception of the dataset by Quan and Ren (2009) (which is problematic in having only *one* annotator per sentence), the largest gold standard for any emotion format (both dimensional and categorical). Even compared with polarity corpora it is still reasonably large (e.g., similar in size to the *Stanford Sentiment Treebank* (Socher et al., 2013)).

4 Analysis and Results

For continuous, real-valued numbers, well-known metrics for IAA, such as Cohen’s κ or F-score, are inappropriate as these are designed for nominally scaled variables. Instead, Pearson’s correlation coefficient (r) or Mean Absolute Error (MAE) are often applied for this setting (Strapparava and Mihalcea, 2007; Yu et al., 2016). Accordingly, for each annotator, we compute r and MAE between their own and the aggregated EMOBANK annotation and average these values for each VAD dimension. This results in one IAA value per metric (r or MAE), perspective and dimension (Table 2).

As average over the VAD dimensions, we achieve a satisfying IAA of $r > .6$ for both perspectives. The READER results in significantly higher correlation,⁵ but also higher error than

⁵Note that using this set-up, obtaining statistical significance is very rare, since the number of cases is based on the number of raters.

	Valence	Arousal	Dominance	Av.
r_{writer}	0.698	0.578	0.540	0.605
r_{reader}	0.738	0.595	0.570	0.634
MAE_{writer}	0.300	0.388	0.316	0.335
MAE_{reader}	0.349	0.441	0.367	0.386

Table 2: IAA for the three VAD dimensions.

WRITER ($p < .05$ for Valence in r and for all dimensions in MAE using a two-tailed t -test).

Prior work found that a large portion of language may actually be neutral in terms of emotion (Ovesdotter Alm et al., 2005). However, a too narrow rating distribution (i.e., most of the ratings being rather neutral relative to the three VAD dimensions) may be a disadvantageous property for training data. Therefore, we regard the *emotional-ity* of ratings as another quality criterion for emotion annotation complementary to IAA.

We capture this notion as the absolute difference of a sentence’s aggregated rating from the neutral rating (3, in our case), averaged over all VAD dimensions. Comparing the average emotionality of all sentences between WRITER and READER, we find that the latter perspective also excels with significantly higher emotionality than the WRITER ($p < .001$; two-tailed t -test).

These beneficial characteristics of the READER perspective (better correlation-based IAA and emotionality) contrast with its worse error-based IAA. Thus, we decided to examine the relationship between error and emotionality between the two perspectives more closely: Let V, A, D be three $m \times n$ -matrices where m corresponds to the number of sentences and n to the number of annotators so that the three matrices yield all the individual ratings for Valence, Arousal and Dominance, respectively. Then we define the *sentence-wise error* for sentence i (SWE_i) as

$$SWE_i := \frac{1}{3} \sum_{X \in \{V, A, D\}} \frac{1}{n} \sum_{j=1}^n |\bar{X}_i - X_{ij}| \quad (1)$$

where $\bar{X}_i := \frac{1}{n} \sum_{j=1}^n X_{ij}$. We compute SWE values for reader and writer perspective individually. We can now examine the dependency between error and emotionality by subtracting, for each sentence, SWE and emotionality for both perspectives from another (resulting in one *difference in error* and one *difference in emotionality* value).

Our data reveal a strong correlation ($r = .718$) between these data series, so that the more the ratings for a sentence differ in emotionality (compar-




Figure 3: Differences in emotionality and differences in error between WRITER and READER, each sentence corresponding to one data point; regression line depicted in red.

ing between the perspectives), the more they differ in error as well. Running linear regression on these two data rows, we find that the regression line runs straight through the origin (intercept is *not* significantly different from 0; $p = .992$; see Figure 3). This means that without difference in emotionality, WRITER and READER rating for a sentence do, on average, *not* differ in error. Hence, our data strongly suggest that READER is the superior perspective yielding better inter-annotator *correlation* and emotionality without overproportionally increasing inter-annotator *error*.

5 Mapping between Emotion Formats

Making use of the bi-representational subset of our corpus (SE07), we now examine the feasibility of automatically mapping between dimensional and categorical models. For each Basic Emotion category, we train one k Nearest Neighbor model given all VAD values of either WRITER, READER or both combined as features. Training and hyperparameter selection was performed using 10-fold cross-validation.

Comparing the correlation between our models' predictions and the actual annotations (in categorical format) with the IAA as reported by Strapparava and Mihalcea (2007), we find that this approach already comes close to human performance (see Table 3). Once again, READER turns out to be superior in terms of the achieved mapping performance compared to WRITER. However, both perspectives combined yield even better results. In this case, our models' correlation with the actual SE07 rating is as good as or even better than the average human agreement. Note that the SE07 ratings are in turn based on averaged human judgments. Also, the human IAA differs a lot between

	Joy	Ang	Sad	Fea	Dsg	Srp	Av.
IAA	.60	.50	.68	.64	.45	.36	.54
W	.68	.40	.67	.47	.27	.15	.44
R	.73	.47	.68	.54	.36	.15	.49
WR	.78	.50	.74	.56	.36	.17	.52
D _W	+.08	-.10	-.01	-.17	-.17	-.21	-.09
D _R	+.13	-.03	+.00	-.10	-.09	-.22	-.05
D _{WR}	+.18	+.00	+.05	-.08	-.09	-.19	-.02

Table 3: IAA by Strapparava and Mihalcea (2007) compared to mapping performance of KNN models using writer's, reader's or both's VAD scores as features (W, R and WR, respectively), both in Pearson's r . Bottom section: difference of respective model performance (W, R and WR) and IAA.

the Basic Emotions and is even $r < .5$ for Disgust and Surprise. For the four categories with a reasonable IAA, Joy, Anger, Sadness and Fear, our best models, on average, actually outperform human agreement. Thus, our data shows that automatically mapping between representation formats is feasible at a performance level on par with or even surpassing human annotation capability. This finding suggests that, for a dataset with high-quality annotations for one emotion format, automatic mappings to another format may be just as good as creating these new annotations by manual rating.

6 Conclusion

We described the creation of EMOBANK, the first large-scale corpus employing the dimensional VAD model of emotion and one of the largest gold standards for *any* emotion format. This genre-balanced corpus is also unique for having two kinds of double annotations. First, we annotated for both writer and reader emotion; second, for a subset of the EMOBANK, ratings for categorical Basic Emotions as well as VAD dimensions are now available. The statistical analysis of our corpus revealed that the reader perspective yields both better IAA values and more emotional ratings. For the bi-representationally annotated subcorpus, we showed that an automatic mapping between categorical and dimensional formats is feasible with near-human performance using standard machine learning techniques.

Acknowledgments

We thank The Center for the Study of Emotion and Attention, University of Florida, for granting us access to the Self-Assessment-Manikin (SAM).

References

- Margaret M. Bradley and Peter J. Lang. 1994. Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry*, 25(1):49–59.
- Margaret M. Bradley and Peter J. Lang. 1999. Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings. Technical Report C-1, The Center for Research in Psychophysiology, University of Florida, Gainesville, FL.
- Margaret M. Bradley and Peter J. Lang. 2007. Affective norms for English text (ANET): Affective ratings of text and instruction manual. Technical Report D-1, The Center for Research in Psychophysiology, University of Florida, Gainesville, FL.
- Sven Buechel and Udo Hahn. 2016. Emotion analysis as a regression problem: Dimensional models and their implications on emotion representation and metrical evaluation. In Gal A. Kaminka, Maria Fox, Paolo Bouquet, Eyke Hüllermeier, Virginia Dignum, Frank Dignum, and Frank van Harmelen, editors, *ECAI 2016 — Proceedings of the 22nd European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS 2016). The Hague, The Netherlands, August 29 - September 2, 2016*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 1114–1122, Amsterdam, Berlin, Washington, D.C. IOS Press.
- Sven Buechel and Udo Hahn. 2017. Readers vs. writers vs. texts: Coping with different perspectives of text understanding in emotion annotation. In *LAW 2017 — Proceedings of the 11th Linguistic Annotation Workshop. Valencia, Spain, April 3, 2017*.
- Rafael A. Calvo and Sunghwan Mac Kim. 2013. Emotions in text: Dimensional and categorical models. *Computational Intelligence*, 29(3):527–543.
- Paul Ekman. 1992. An argument for basic emotions. *Cognition & Emotion*, 6(3-4):169–200.
- Nancy C. Ide, Collin F. Baker, Christiane Fellbaum, Charles J. Fillmore, and Rebecca J. Passonneau. 2008. MASC: The Manually Annotated Sub-Corpus of American English. In Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan E. J. M. Odijk, Stelios Piperidis, and Daniel Tapias, editors, *LREC 2008 — Proceedings of the 6th International Conference on Language Resources and Evaluation. Marrakech, Morocco, 26 May - June 1, 2008*, pages 2455–2461.
- Nancy C. Ide, Collin F. Baker, Christiane Fellbaum, and Rebecca J. Passonneau. 2010. The Manually Annotated Sub-Corpus: A community resource for and by the people. In Jan Hajič, M. Sandra Carberry, and Stephen Clark, editors, *ACL 2010 — Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics. Uppsala, Sweden, 11-16 July 2010*, volume 2: Short Papers, pages 68–73.
- Phil Katz, Matthew Singleton, and Richard Wicentowski. 2007. SWAT-MP: The SemEval-2007 systems for Task 5 and Task 14. In Eneko Agirre, Lluís Màrquez, and Richard Wicentowski, editors, *SemEval-2007 — Proceedings of the 4th International Workshop on Semantic Evaluations @ ACL 2007. Prague, Czech Republic, June 23-24, 2007*, pages 308–313.
- Maximilian Köper and Sabine Schulte im Walde. 2016. Automatically generated affective norms of abstractness, arousal, imageability and valence for 350,000 German lemmas. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asunción Moreno, Jan E. J. M. Odijk, and Stelios Piperidis, editors, *LREC 2016 — Proceedings of the 10th International Conference on Language Resources and Evaluation. Portorož, Slovenia, 23-28 May 2016*, pages 2595–2598.
- Peter J. Lang. 1980. Behavioral treatment and bio-behavioral assessment: Computer applications. In J. B. Sidowski, J. H. Johnson, and T. A. Williams, editors, *Technology in Mental Health Care Delivery Systems*, pages 119–137. Ablex, Norwood/NJ.
- Shoushan Li, Jian Xu, Dong Zhang, and Guodong Zhou. 2016. Two-view label propagation to semi-supervised reader emotion classification. In Nicoletta Calzolari, Yuji Matsumoto, and Rashmi Prasad, editors, *COLING 2016 — Proceedings of the 26th International Conference on Computational Linguistics. Osaka, Japan, December 11-16, 2016*, volume Technical Papers, pages 2647–2655.
- Huanhuan Liu, Shoushan Li, Guodong Zhou, Chu-Ren Huang, and Peifeng Li. 2013. Joint modeling of news reader’s and comment writer’s emotions. In Hinrich Schütze, Pascale Fung, and Massimo Poesio, editors, *ACL 2013 — Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics. Sofia, Bulgaria, August 4-9, 2013*, volume 2: Short Papers, pages 511–515.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Agnes Moors, Jan De Houwer, Dirk Hermans, Sabine Wanmaker, Kevin van Schie, Anne-Laura Van Harmelen, Maarten De Schryver, Jeffrey De Winne, and Marc Brysbaert. 2013. Norms of valence, arousal, dominance, and age of acquisition for 4,300 Dutch words. *Behavior Research Methods*, 45(1):169–177.
- Cecilia Ovesdotter Alm, Dan Roth, and Richard Sproat. 2005. Emotions from text: Machine learning for text-based emotion prediction. In Raymond J. Mooney, Christopher Brew, Lee-Feng Chien, and Katrin Kirchhoff, editors, *HLT-EMNLP 2005 — Proceedings of the Human Language Technology Conference & 2005 Conference on Empirical*

- Methods in Natural Language Processing. Vancouver, British Columbia, Canada, 6-8 October 2005*, pages 579–586.
- Robert Plutchik. 1980. A general psychoevolutionary theory of emotion. *Emotion: Theory, Research and Experience*, 1(3):3–33.
- Daniel Preoțiuc-Pietro, Hansen Andrew Schwartz, Gregory Park, Johannes C. Eichstaedt, Margaret L. Kern, Lyle H. Ungar, and Elizabeth P. Shulman. 2016. Modelling valence and arousal in Facebook posts. In Alexandra Balahur, Erik van der Goot, Piek Vossen, and Andrés Montoyo, editors, *WASSA 2016 — Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis @ NAACL-HLT 2016. San Diego, California, USA, June 16, 2016*, pages 9–15.
- Changqin Quan and Fuji Ren. 2009. Construction of a blog emotion corpus for Chinese emotional expression analysis. In Philipp Koehn and Rada Mihalcea, editors, *EMNLP 2009 — Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing. A Meeting of SIGDAT, a Special Interest Group of ACL @ ACL-IJCNLP 2009. Singapore, 6-7 August 2009*, pages 1446–1454.
- James A. Russell and Albert Mehrabian. 1977. Evidence for a three-factor theory of emotions. *Journal of Research in Personality*, 11(3):273–294.
- David Sander and Klaus R. Scherer, editors. 2009. *The Oxford Companion to Emotion and the Affective Sciences*. Oxford University Press, Oxford; New York.
- David S. Schmidtke, Tobias Schröder, Arthur M. Jacobs, and Markus Conrad. 2014. ANGST: Affective norms for German sentiment terms, derived from the affective norms for English words. *Behavior Research Methods*, 46(4):1108–1118.
- Kim Schouten and Flavius Frasinca. 2016. Survey on aspect-level sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering*, 28(3):813–830.
- Yohei Seki, Noriko Kando, and Masaki Aono. 2009. Multilingual opinion holder identification using author and authority viewpoints. *Information Processing & Management*, 45(2):189–199.
- Parinaz Sobhani, Saif M. Mohammad, and Svetlana Kiritchenko. 2016. Detecting stance in tweets and analyzing its interaction with sentiment. In Claire Gardent, Raffaella Bernardi, and Ivan Titov, editors, **SEM 2016 — Proceedings of the 5th Joint Conference on Lexical and Computational Semantics @ ACL 2016. Berlin, Germany, August 11-12, 2016*, pages 159–169.
- 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. In Timothy Baldwin and Anna Korhonen, editors, *EMNLP 2013 — Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing. Seattle, Washington, USA, 18-21 October 2013*, pages 1631–1642.
- Hans Stadthagen-Gonzalez, Constance Imbault, Miguel A. Pérez Sánchez, and Marc Brysbaert. 2016. Norms of valence and arousal for 14,031 Spanish words. *Behavior Research Methods*. 10.3758/s13428-015-0700-2.
- Carlo Strapparava and Rada Mihalcea. 2007. SemEval-2007 Task 14: Affective text. In Eneko Agirre, Lluís Màrquez, and Richard Wicentowski, editors, *SemEval-2007 — Proceedings of the 4th International Workshop on Semantic Evaluations @ ACL 2007. Prague, Czech Republic, June 23-24, 2007*, pages 70–74.
- Carlo Strapparava. 2016. Emotions and NLP: Future directions. In Alexandra Balahur, Erik van der Goot, Piek Vossen, and Andrés Montoyo, editors, *WASSA 2016 — Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis @ NAACL-HLT 2016. San Diego, California, USA, June 16, 2016*, page 180.
- Yi-jie Tang and Hsin-Hsi Chen. 2012. Mining sentiment words from microblogs for predicting writer-reader emotion transition. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Asunción Moreno, Jan E. J. M. Odijk, and Stelios Piperidis, editors, *LREC 2012 — Proceedings of the 8th International Conference on Language Resources and Evaluation. Istanbul, Turkey, May 21-27, 2012*, pages 1226–1229.
- Jin Wang, Liang-Chih Yu, K. Robert Lai, and Xuejie Zhang. 2016. Dimensional sentiment analysis using a regional CNN-LSTM model. In Antal van den Bosch, Katrin Erk, and Noah A. Smith, editors, *ACL 2016 — Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics. Berlin, Germany, August 7-12, 2016*, volume 2: Short Papers, pages 225–230.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4):1191–1207.
- Janyce M. Wiebe, Theresa Ann Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3 (Special Issue on “Advances in Question Answering”)):165–210.
- Liang-Chih Yu, Jin Wang, K. Robert Lai, and Xuejie Zhang. 2015. Predicting valence-arousal ratings of words using a weighted graph method. In Yuji Matsumoto, Chengqing Zong, and Michael Strube, editors, *ACL-IJCNLP 2015 — Proceedings of the 53rd*

Annual Meeting of the Association for Computational Linguistics & 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing. Beijing, China, July 26-31, 2015, volume 2: Short Papers, pages 788–793.

Liang-Chih Yu, Lung-Hao Lee, Shuai Hao, Jin Wang, Yunchao He, Jun Hu, K. Robert Lai, and Xuejie Zhang. 2016. Building Chinese affective resources in valence-arousal dimensions. In Kevin C. Knight, Ani Nenkova, and Owen Rambow, editors, *NAACL-HLT 2016 — Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. San Diego, California, USA, June 12-17, 2016*, pages 540–545.