Annotated Bibliography

1. Alm, C. O. (2012). The role of affect in the computational modeling of *natural* language. *Language and Linguistics Compass (Computational and Mathematical)*, 6(7), 416-430. SURVEY (GENERAL).
Written for interdisciplinary readers across career-levels, this succinct survey article provides literature review and discussion about theoretical background and applied topics of interest for analyzing affect in linguistic corpora and for incorporating mechanisms for processing affect into language technology.

2. Cahn, J. (1990). The generation of affect in synthesized speech. *Journal of the American Voice I/O Society*, 8, 1-19. SEMINAL/CLASSICAL.
This classical reading represents an early example of work on synthesizing emotional speech. It describes the Affect Editor, a tool for generating expressive speech, including a diverse set of speech parameters used for modeling, and it discusses the implementation and results of human-based evaluation. The work suggested that recognizable emotions could be synthesized. Some of the findings included that sad stimuli were particularly well identified, that lexical contents of sentences might influence perception, and that category mix-up might occur between more affectively similar concepts such as angry and disgusted. The paper also briefly reported on observing individual preferences. Affective expressiveness continues to be a research topic in speech science and technology communities.

3. Calvo, R. A., and D'Mello, S. (2010). Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing*, 1(1), 18-37. SURVEY (GENERAL).
This is a comprehensive survey article useful for readers who wish to deepen their understanding of the interdisciplinary landscape of work involved with detecting affect, the affective sciences, and affective computing. Given theoretical background and frameworks for modeling affect, the article dives into an overview of studied human modalities that contribute affect signals (with particular sections dedicated to spoken and written language), including methods, resources, and multimodal integration. Discussion synthesizes important topics and current/future research directions. The article presents an opportunity to become familiar with IEEE Transactions on Affective Computing.

4. Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. (2011). Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12, 2493–2537. EMPIRICAL (SYSTEM).
This paper presents a neural network framework that was applied to part-of-speech tagging, chunking, named entity recognition, semantic role labeling, and some other NLP tasks. It is one of the more recent papers on deep learning in NLP that eschews task-specific engineering in favour of learning common internal representations from data. Even though affect-related tasks are not directly addressed in this paper, several deep learning papers on valence classification draw inspiration from this work.

5. Coppersmith, G., Dredze, M., and Harman, C. (2014). Quantifying mental health signals in Twitter. *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality at the 52nd Annual Meeting of the Association for Computational Linguistics, Baltimore, MD, USA*, 51-60. APPLICATION (HEALTH).
This study from a recent workshop exemplifies the interest in textual data and social media for exploring affect-related phenomena in the domain of public health. Presenting an approach for gathering and studying microblog data for conditions such as depression and PTSD, the paper also discusses the opportunity for complementary use of natural language processing techniques in relation to a tradition of survey analysis in health contexts. The paper conveys some of the challenges (generational uses of social media, less mention of uncommon conditions, etc.) as well as the usefulness of interdisciplinary collaboration in this area.

6. Cornelius, R. R. (2000). Theoretical approaches to emotion. *Proceedings of the ISCA ITRW on Speech and Emotion (SpeechEmotion-2000)*, Newcastle, Northern Ireland, UK, 3-10. SURVEY (THEORY).
Cornelius advocates for the need to take theoretical accounts into consideration in emotion scholarship. He straightforwardly introduces four “perspectives” from the discipline of psychology: Darwinian, Jamesian, cognitive, and social constructivist, and also discusses how these views relate to each other. While positioning the discussion within emotional speech research, the author explicates some of the benefits of understanding where one’s work fits theoretically and how theoretical views are merging, including how that may aid the appreciation of assumptions involved or influence investigatory questions and insights to evolve.

7. Cowie, R., Douglas-Cowie, E., Martin, J.-C., and Devillers, L. (2010). The essential role of human databases for learning in and validation of affectively competent agents. In Scherer, K. R., Bänziger, T., and Roesch, E. B. (Eds.) *Blueprint for Affective Computing: A Sourcebook*. Oxford: Oxford University Press, 151-165. SURVEY (RESOURCE).
This book chapter about data resources touches upon many of the issues and topics under discussion with respect to affect data development, from the perspective of the affect sciences and affective computing. Database examples are covered for distinct modalities (albeit sparsely for text-oriented work). The chapter includes summarizing projections about next developments in this area.

8. Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., Weeg, C., Larson, E. E., Ungar, L. H., and Seligman, M. E. (2015). Psychological language on Twitter predicts county-level heart disease mortality. *Psychological Science,* 26(2), 159-169. APPLICATION (HEALTH).
Traditional approaches for determining psychological states of people involve in-person or phone conversations. Such approaches are time intensive and expensive. This paper aims at determining psychological state, at the

community level, from tweets posted by the community. Specifically, it analyzes language in tweets and finds correlations of certain features with heart disease rates at the level of counties. Lexicons for anger, anxiety, positive and negative emotions, positive and negative social relationships, and engagement and disengagement were found to be useful. This is an interesting example of bridging the analysis of language (in this case tweets) with non-linguistic information (in this case heart disease data).

9. Kiritchenko, S., Zhu, X., and Mohammad, S. M. (2014). Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 50, 723–762. EMPIRICAL (SYSTEM).
This paper gives details about the NRC-Canada system that came first in various sentiment-related shared tasks in Sem-Eval-2013 and 2014. A number of different kinds of features were used. Ablation experiments showed that the Twitter-specific valence-association lexicons were the most useful. These automatically generated lexicons capture non-standard language, such as creative spellings and word elongations. Methods to capture impact of negation on sentiment are also described. The paper additionally describes a maximum difference scaling approach to obtain reliable fine-grained annotations of sentiment.

10. Liu, H., Lieberman, H., and Selker, T. (2003). A model of textual affect sensing using real-world knowledge. *Proceedings of the International Conference on Intelligent User Interfaces, Miami, FL, USA,* 125-132. SEMINAL/CLASSICAL.
In this early paper on affect processing with text, the main application of interest was an affective email interface (EmpathyBuddy). The work involved interpreting affect in terms of fundamental emotion categories, using a textual resource of commonsense knowledge (Open Mind Commonsense). A user study evaluated the email client. Users interacted with three client versions, including the sensing-based version. They assessed the system for "entertainment, interactivity, intelligence, and adoption". The results suggested that the authors' main approach was perceived as more intelligent, adoptable, and interactive.

11. Mohammad, S. M., and Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29 (3), 436-465. EMPIRICAL (RESOURCE).
This paper describes the creation of a large word—affect association lexicon by crowdsourcing. Several techniques are employed for quality control, most notably with the use of a separate word-choice question that ascertains whether the annotator knows the meaning of the target word. The question also guides the annotator to the desired sense of the word for which annotations are needed. The resulting lexicon, the NRC Emotion Lexicon, has entries for over 14,000 words and about 25,000 word senses. Each instance is marked for associations with eight emotions, as well as positive and negative sentiment. The lexicon is widely used by researchers and system builders for various affect-related tasks.

12. Mohammad, S. M., and Kiritchenko, S. (2013). Using nuances of emotion to identify personality. *Proceedings of the International Conference on Weblogs and Social Media (ICWSM-13)*, *Boston, MA*, 27-30. EMPIRICAL (SYSTEM).
This paper describes the collection and use of tweets with emotion word hashtags for automatic emotion detection. Experiments show that the emotion word hashtags act as good labels of emotions in the rest of the tweet. Thus the data can be used for training machine learning systems for emotion classification. The success of this approach also means that one can now quickly compile training data for any emotion which is used as a hashtag in tweets. Experiments are performed in an extrinsic task for personality trait classification, where it is shown that emotion-based features from hundreds of emotion categories are more useful than using features from a handful of affect categories (such as positive and negative sentiment, or the Big Six emotion categories).

13. Neviarouskaya, A., Prendinger, H., and Ishizuka, M. (2010). Recognition of affect, judgment, and appreciation in text. *Proceedings of the 23rd International Conference on Computational Linguistics, Beijing, China*, 806-814. EMPIRICAL (SYSTEM).
Even though a majority of current approaches in NLP are statistical, rule-based systems are more interpretable, for example, for understanding why a sentence was classified as having a certain emotion by the system. This paper presents a rule-based system for detecting emotions at sentence level. It employs a number of manually created lexicons for attitude, affect modifiers, and even a lexicon that captures the confidence signified by modal verbs. At the heart of the system is a method to combine affect-related information from different pieces of the text using rules. Developing compositional models statistically is a major area of research these days, and this work can be a source of ideas in developing composition models that capture affect appropriately.

14. Ortony, A., Clore, G. L., and Collins, A. (1990). *The Cognitive Structure of Emotions*. Cambridge: Cambridge University Press. SEMINAL/CLASSICAL.
Several authors have proposed mutually conflicting theories about emotions, and till date many key aspects of emotions are hotly debated. This book by Ortony, Clore, and Collins presents one such theoretical framework that argues that emotions are valence reactions. The valence reaction is broken down into several sub-categories and these subcategories are broken down into further sub-categories, based on whether the valence reaction was to the consequences of events, aspects of objects, whether the person approves or disapproves it, etc. Ideas on the theoretical underpinnings of emotions and different kinds of valence reaction can be helpful for developing instructions for affect annotations, as well as for developing features that can be useful in automatic affect classification.

15. Osgood, C. E., Suci, G. J., and Tannenbaum, P. (1957). *The Measurement of Meaning*. Urbana, USA: University of Illinois Press. SEMINAL/CLASSICAL.
This work studies the nature of meaning. One of its most influential experiments involves asking people to rate the meanings of concepts along several dimensions such as fair—unfair, strong—weak, safe—dangerous, etc. Factor analysis of the responses is used to show that the three dimensions conveying most of the variance in meaning across concepts are that of evaluativeness (good—bad), potency (strong—weak), and activity (active—passive). This work has influenced scholars in a number of fields including linguistics, psychology, mass communications, and natural language processing.

16-17. Brief mention of two book manuscripts. SURVEY (GENERAL).
Scherer, K. R., Bänziger, T., and Roesch, E. B. (Eds.) (2010). *Blueprint for Affective Computing: A Sourcebook*. Oxford: Oxford University Press.
In this collection, the editors gather nineteen chapters into seven sections on a range of topics pertinent for exploring affect in human- and machine-oriented research. Chapters range from topics such as "Emotions in interpersonal interactions" (Parkinson) to "Emotion in artificial neural networks" (Roesch, Korsten, Fragopanagos, and Taylor), with five chapters specifically devoted to "Approaches to an implementation of affectively competent agents".
Schuller, B., and Batliner, A. (2014). *Computational Paralinguistics: Emotion, Affect and Personality in Speech and Language Processing*. Chichester: Wiley.
This recent book is divided into "Foundations" and "Modelling", each comprising various chapters, such as "Taxonomies", "Functional aspects", and "Corpus engineering" in the first part, and "Linguistic features", "Machine-based modelling", and "'Hands-on': Existing toolkits and practical tutorial" in the second part. In the preface, the authors explain that a goal is "to provide the reader with a sort of map presenting an overview of the field, and useful for finding one's way through. The scale of this map is medium-sized, and we can only display a few of the houses in this virtual paralinguistic 'city' with their interiors, on an exemplary basis."

18. Strapparava, C., and Mihalcea, R. (2007). SemEval-2007 Task 14: Affective text. *Proceedings of SemEval-2007*, *Prague, Czech Republic,* 70–74. EMPIRICAL (RESOURCE).
This is a task-description paper of an early shared task competition on automatically detecting valence and emotions in text. The dataset chosen was a collection of newspaper headlines. Annotators were asked to give scores between 0 and 100 for each of the Big Six emotions and positive and negative valence. No training data was provided, and so only unsupervised systems were able to participate. Nonetheless, the testset created as part of this shared task has subsequently been used by supervised systems by splitting it into new train—test partitions. One of the key distinctions of this work compared to the shared tasks proposed in the last few years is that this is one of the few datasets that has fine-grained annotation for the degree of

affect. Several applications would benefit from a system that can predict the degree of affect in text.

19. Turney, P.D. (2002). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, Philadelphia, PA*, 417-424. EMPIRICAL (SYSTEM).
This paper presents a way to classify customer reviews as positive (recommended) or negative (not recommended). At the heart of the method is a way to determine the degree of positiveness (or negativeness) of a word by calculating the mutual information score between it and sets of positive and negative seed terms. This fundamental approach is still used (possibly with minor modification) for creating sentiment and emotion association lexicons.

20. Wiebe, J., Wilson, T., Bruce, R., Bell, M., and Martin, M. (2004). Learning subjective language. *Computational Linguistics*, 30(3), 277–308. SEMINAL/CLASSICAL.
This article presents early, comprehensive efforts to automatically detect subjective language. Supervised classification is performed on a number of datasets using various features, including low-frequency words, collocations, and words that are distributionally similar to pre-chosen seed words. The paper is also an excellent resource for understanding the principles underpinning subjective language.