The Workshops of the Thirty-Second
AAAI Conference on Artificial Intelligence

The AAAI-18 Workshop on Affective Content Analysis

Niyati Chhaya
Big Data Experience Lab
Adobe Research, India
nchhaya@adobe.com

Kokil Jaidka
School of Arts and Sciences
University of Pennsylvania, USA
jaidka@sas.upenn.edu

Lyle H. Ungar
Department of Computer
and Information Science
University of Pennsylvania, USA
ungar@cis.upenn.edu

Abstract

The first AAAI-18 Workshop on Affective Content Analysis was an interdisciplinary platform that focused on the analysis of emotions, sentiments, and attitudes in textual, visual, and multimodal content for applications in psychology, consumer behavior, language understanding, and computer vision. The program comprised interdisciplinary keynotes, original research presentations, a poster session and short pitches for datasets and pre-published work.

Introduction

The effectiveness of any published material is driven by the text/content stated (by the author) as well as the reaction it achieves (from the reader). This content ranges from facts to opinions, reports to advertisements and marketing material, and social media to personal emails. The reader who reacts to the published or authored material interprets the content and reacts based on their perception. This perception evolves from various different aspects of the reader's personality; an important dimension being the psycho–demographic and psychographic orientation of the reader. Psychologists and Consumer behavior researchers work on building theories that aim to quantify such human reactions.

Affective Computing has traditionally focused on modeling human reactions using multi–modal sensor data (Picard and Picard 1997). This community, however, has focused mainly on non–text data. Sentiment and emotion analysis on the other hand has been applied on text as well as multi–modal datasets, but this research has been limited to quantifying well–defined human reactions. The affect analysis i.e. techniques and applications that understand the 'experience of an emotion' (Picard and Picard 1997) in the context of language and text is a upcoming research space.

The goal of the AAAI Workshop on Affective Content Analysis is to provide a platform to stimulate interdisciplinary discussions for affect in content, and engage the AI and ML community about the open problems in affective content analysis and understanding, with a special focus on affect in language and text. Affective content analysis in this context refers to the interdisciplinary research space of computational linguistics, psycholinguists, consumer psychology, and HCI looking at online communication in various forms. Work on affect analysis in language and text spans many research communities, including computational linguistics, consumer psychology, human-computer interaction (HCI), marketing science and cognitive science. Computational linguists study how language evokes as well as expresses emotion. Consumer psychology examines human affect by drawing upon grounded psychological theories of human behavior. The HCI community studies human responses as a part of user experience evaluation. Computational models for consumer psychology theories present a huge opportunity to guide the construction of intelligent systems that understand human reactions, and tools from linguistics and machine learning can provide attractive methods to fulfill those opportunities. Models of affect have recently been adapted for social media platforms, enabling new approaches to understanding users' opinions, intentions and expressions.

The exponentially increasing size and the dynamic, multimedia nature of this data make it difficult to detect and measure affect. Furthermore, the subjective nature of human affect suggests the need to measure in ways that recognize multiple interpretations of human responses. Other key challenges in this space are:

- Standardizing the measurement of affect in order to meaningfully compare different affective models against each other
- Cross-media, cross-domain and cross-platform affect analysis
- Building a theoretical framework of affect based on the literature in consumer psychology and cognition
- Building language-based affect models as an input for other data science applications
- Building annotated datasets, standardized baselines and developing metrics for meaningful evaluation and benchmarking.

The AI community provides the appropriate middle ground for bringing together researchers from multiple disciplines for stimulating discussions on the open research problems in affect analysis.

Workshop Topics and Format

Papers and talks at the workshop incorporated insights from psychologists, consumer behavior researchers, and computational linguistics to develop new approaches that address open problems such as deep learning for affect analysis, leveraging traditional affective computing algorithms (multi–modal data and sensors) for text, measurement of affect and its modeling, and understanding and identifying the right affect–related dimensions to study consumer behavior. These fall under the umbrella of topics of interest to the workshop:

- Affect and Cognitive Content Measurement in Text
- Computational models for Consumer Behavior theories
- Psycho-demographic Profiling
- Affect–based Text Generation
- Spoken and Formal Language Comparison
- Stylometrics, Typographics, and Psycho-linguistics
- Affective needs and Consumer Behavior
- Measurement and Evaluation of Affective Content
- Affective Lexica for Online Marketing Communication
- Affective Commonsense Reasoning
- Affective human-agent, -computer, and-robot interaction
- Multi-modal emotion recognition and sentiment analysis

Given the need of standardized baselines, datasets, and evaluation metrics, the workshop also had a session dedicated to the datasets and resources available for multimodal affect analysis in the different domains.

Overview of the papers

The workshop featured five keynote talks, two paper sessions, and a poster session. 29 papers were submitted to the workshop, 7 of which were pre-published works. Finally, 5 papers were accepted as full papers and 1 was accepted as a short paper for inclusion in the proceedings. In addition, 7 papers were invited for the poster session. The following sections briefly describe the keynote and sessions.

Keynote

The workshop had a range of keynote speakers. Dr. Dipankar Chakravarti [1] and Dr. Rajesh Bagchi [2] shared their work in the space of consumer psychology and marketing science. Dr. Bagchi discussed the theoretical and managerial implications of his work with consumers and calorie consumption. Over a discussion of three studies, he demonstrated that when consumers process information affectively, they consume more calories, because of the relationships between distraction, cognitive thinkings, and affective processing.

Dr. James Pennebaker, a well known social psychologist [3] talked about affective language understanding. His talk provided evidence from a series of studies about how the function words people use, such as pronouns, prepositions, and other common parts of speech, it is possible to detect the ways in which people pay attention to their social world across different emotional states.

Dr. Cristian Danescu-Niculescu-Mizil [4] talked about conversational dynamics and their relations to social interactions. In the talk a computational framework for modeling conversations dynamics and social signal encoding was presented. In particular, temporal friendships and betrayal are characterized using Diplomacy strategy game. Dr.Danescu-Niculescu-Mizil also discussed his findings from the study of group discussions, especially the effect of over– and under–confidence on dynamics and outcomes of decision-making discussions.

Last but not the least, Dr. Jennifer Healey [5] discussed her work in multimodal affect analysis and the possibilities of a promising future, with devices and systems that are sensitive to human emotions.

Session 1: Affect in Text

The paper, ‘Why is an Event Affective? Classifying Affective Events based on Human Needs’ (Ding, Jiang, and Riloff 2018) introduces a classification schema for categorizing affective events depending on human needs they relate to. It further presents a frame-like even structure used for extracting and representing events and describes experiments conducted on the gold standard dataset. The applied baseline classification methods: rule based system using LIWC and supervised classifiers over ngram and event embedding features achieve a moderate performance.

The paper titled ‘Emotion Detection on TV Show Transcripts with Sequence-based Convolutional Neural Networks’ (Zahiri and Choi 2018) introduces a new data set of transcriptions of dialogues from a TV show annotated with 6 emotions and neutral and proposes a number of CNN based models that can predict emotions by taking into account previous states and texts. The novelty of this work lies in the modeling the entire sequence of dialogue in predicting emotion. The authors report an accuracy of 37.9% and 54% for fine- and coarse-grained emotions, respectively.

The paper titled ‘Do convolutional networks need to be deep for Text Classification?’ (Le, Cerisara, and Denis 2018) also applies convolutional models to understand the effect of the depth of a deep convolution model for text and sentiment classification. The authors present a comprehensive set of experiments with shallow-and-wide networks and dense-net networks, for 5 different datasets with different classification objective, and compare their findings against 13 baselines. Unexpectedly, their shallow-and-wide networks outperform deep models when inputs are words, but deep models outperform shallow networks when inputs are a sequence of characters. The findings have important implications for

[1] https://marketing.pamplin.vt.edu/people/faculty/chakravarti-dipankar.html

[2] https://marketing.pamplin.vt.edu/people/faculty/bagchi-rajesh.html

[3] https://liberalarts.utexas.edu/psychology/faculty/pennebak

[4] www.cs.cornell.edu/ cristian/

[5] https://www.linkedin.com/in/healeyjennifer

how future sentiment analysis tasks should be modeled, and set a new benchmark performance for the Yelp datasets using shallow word models.

The paper 'Knowledge driven feed-forward neural network for audio affective content analysis' (Dumpala, Chakraborty, and Kopparapu 2018) improves Feed-Forward Neural Network (FFNN) performance on induced emotions from the LIRIS-ACCEDE database. The authors propose the novel idea of incorporating prior knowledge into training, by applying a weighting function the target outputs during training. This downweights the importance of samples at the beginning and end of a segment. The authors demonstrate that their method outperforms standard feed-forward neural networks and recurrent neural networks on the MediaEval dataset, especially when the training data is sparse.

Session 2: Multimodal affect

In their paper titled 'Predicting Engagement Breakdown in HRI Using Thin-slices of Facial Expressions' (Liu and Kappas 2018), the authors present a method for recognizing engagement breakdowns in human-robot interaction from facial expression. Facial action units were detected with Emotionet as features and an echo state recurrent neural networks was trained for the task. The engagement breakdown could be detected withe F1 score of 0.76 which shows the feasibility of the proposed approach. The authors apply their architecture on a real-world dataset to show accurate prediction of engagement breakdown using 30 seconds of facial expressions.

The paper titled 'Multimodal Alignment for Affective Content' (Nester et al. 2018) describes methods for the multimodal alignment of human gaze, semantic description and affect for image/video stimuli. Multiple experiments are conducted to understand the alignment of language use and facial expressions. An interesting problem of image region annotation is also studied by using both modalities and further by selecting only frequent words. The authors evaluate their framework by exploring whether image valence and word frequency filtering impacts alignment results. The paper concludes with a discussion of its applications in image understanding, media accessibility, and multimodal data fusion.

Datasets

Among the dataset papers, the presentation titled 'RusNeuroPsych: Corpus for Study Relations Between Author Demographic, Personality Traits, Lateral Preferences and Affect in Text ' (Litvinova et al. 2018) provided a manually collected corpus of letters to a friend and narratives comprising informal writing describing emotions and opinions in the Russian langauge. The corpus is annotated with information about the authors' gender, age, psychological testing scores and brain laterality preferences, and is freely available on the RusProfiling Lab webpage.

The DesireDB is introduced in the dataset presentation titled 'Modelling Protagonist Goals and Desires in First-Person Narrative' (Rahimtoroghi et al. 2017), which comprises annotations identifying statements of desire, textual evidence for desire fulfillment, and annotations for whether the stated desire is fulfilled given the evidence in the narrative context. The authors demonstrate that their LSTM Skip-Thought model achieves an F-measure of 0.7 on tracking desire fulfillment in this corpus.

The Echo dataset described in the presentation titled 'Linguistic Reflexes of Well-Being and Happiness in Echo' (Wu et al. 2017) comprises a corpus of private micro-blogs from a well-being application called Echo, where users label each written post about daily events with a happiness score between 1 and 9. The authors explore the extent to which different theoretical accounts can explain the variance in the happiness scores, and suggest that recurrence of 'obligation' and 'incompetence' which affect individual well-being are not well-captured in the current set of lexical and semantic resources. The findings highlight an important research gap from both, a natural language processing and a psycholinguistic perspective.

The presentation titled 'Discriminating Between Truthfulness and Deception Using Infrared Thermal Imaging and Peripheral Physiology(Derakhshan, Mikaeili, and Gedeon 2018) presents an interesting new dataset of thermal and physiological data related to deceptive speech, from 32 subjects in two different settings. Their findings report classification accuracies of 61.1% and 60% for thermal signals in the Best Friend and Mock Crime scenarios, respectively.

Posters

Amongst the posters, 'Storytelling Agents with Personality and Adaptivity (Hu et al. 2015)', presents an interesting work that explores the expression of personality and adaptivity through the gestures of personality-aware virtual agents in a storytelling task. The authors conduct a set of experiments that manipulate the agent personality, and report that humans were able to perceive the intended variation in extraversion between different virtual agents.

'Social media and loneliness: Why an Instagram picture may be worth more than a thousand Twitter words' (Pittman and Reich 2016) investigates the connection between social media and the psychological attribute of loneliness, in particular the advantage of image-based social platforms like Instagram over inherently text-based platforms like Twitter. They empirically find that text-based platforms offer little intimacy and are less effective in combating loneliness than image-based platforms. Their quantitative study reports that loneliness may decrease, while happiness and satisfaction with life may increase, as a function of image-based social media use. In contrast, text-based media use appears ineffectual.

The work on 'Learning Lexico-Functional Patterns for First-Person Affect (Reed et al. 2017)' presents a method to learn proxies for these functions from first person narratives using lexico–functional patterns. The authors have created a fine-grained test set for task-based evaluation. They also show that their approach significantly improves prediction accuracies on standard datasets.

Related Workshops

There is a growing number of workshops and conferences related to affective computing which points to the importance of the research problem at hand, as well as the timeliness of this workshop for the AI community. The following workshops focused mainly on text analysis, sentiment, and subjectivity of the text content:

- SENTIRE series: The workshop on Sentiment Elicitation from Natural Text for Information Retrieval and Extraction has been a continuing series for the past few years at ICDM [6]. The organizers of this workshop series are part of the program committee for the proposed workshop.
- WASSA: The workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis is a workshop series that concentrates on sentiment analysis in text and looks at various aspect–based and subjectivity analysis of text in that context. The workshop has been a popular workshop at top NLP conferences such as EMNLP, ACL, and NAACL in recent years [7]. The organizers of this workshop series as well are a part of the program committee of this proposed workshop.

The following workshops focused on the multi-modal, sensory data in their analysis. Text and language analysis is however not the focus of these workshops. This makes the AAAI Workshop on Affective Content Analysis rather unique in its pitch to bring the two communities together.

- The first workshop on Affective Computing (IJCAI 2017) concentrates on measuring human affects based on sensors and wearable devices.
- 1st Workshop on Tools and Algorithms for Mental Health and Wellbeing, Pain, and Distress (MHWPD)
- Multimodal Emotion Recognition Challenge (MEC 2017) @ 2018 Asian Conference on Affective Computing and Intelligent Interaction (AACII)

Other current relevant events include ACII[8], HUMANAIZE[9], and NLP+CSS[10].

Outlook

This workshop received a promising number of submissions and generated a lot of interest among scholars and the industry. The call for datasets was also successful at identifying a number of interesting resources with text, sensor and visual data for affect analysis. The programme comprising interdisciplinary keynotes, original research presentations, a poster session and short pitches for datasets and pre-published work has proven to be a successful and agile format, and we will continue it in the future. We will continue this multi–disciplinary workshop in an attempt to establish the space of computational approaches for affective content analysis.

Acknowledgments

We would like to thank Adobe Research, for their generous funding which made this workshop possible. We thank our program committee members who did an excellent job of reviewing the submissions. All PC members are documented on the AffCon-18 website[11].

[6]http://sentic.net/sentire/

[7]http://optima.jrc.it/wassa2017/

[8]http://acii2017.org/

[9]http://st.sigchi.org/publications/toc/humanize-2017.html

[10]https://sites.google.com/site/nlpandcss/nlp-css-at-acl-2017

[11]https://sites.google.com/view/affcon18/home

References

Berry, M. W.; Browne, M.; and Signer, B. 2001. Topic annotated enron email data set. *Philadelphia: Linguistic Data Consortium*.

Bestgen, Y., and Vincze, N. 2012. Checking and bootstrapping lexical norms by means of word similarity indexes. *Behavior Research Methods* 44(4):998–1006.

Bestgen, Y. 2008. Building affective lexicons from specific corpora for automatic sentiment analysis. In *LREC*.

Bird, S. 2006. Nltk: The natural language toolkit. In *Proceedings of the COLING/ACL on Interactive presentation sessions*, 69–72. Association for Computational Linguistics.

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3(Jan):993–1022.

Blitzer, J.; Dredze, M.; Pereira, F.; et al. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*, volume 7, 440–447.

Bradley, M. M., and Lang, P. J. 1999. Affective norms for english words (anew): Instruction manual and affective ratings. Technical report, Technical report C-1, the center for research in psychophysiology, University of Florida.

Brett, D., and Pinna, A. 2013. The distribution of affective words in a corpus of newspaper articles. *Procedia-Social and Behavioral Sciences* 95:621–629.

Cambria, E.; Fu, J.; Bisio, F.; and Poria, S. 2015. Affectivespace 2: Enabling affective intuition for concept-level sentiment analysis. In *AAAI*, 508–514.

Chikersal, P.; Poria, S.; Cambria, E.; Gelbukh, A.; and Siong, C. E. 2015. Modelling public sentiment in twitter: using linguistic patterns to enhance supervised learning. In *International Conference on Intelligent Text Processing and Computational Linguistics*, 49–65. Springer.

Cohen, W. W. 2009. Enron email dataset.

Colombetti, G. 2009. From affect programs to dynamical discrete emotions. *Philosophical Psychology* 22(4):407–425.

Danescu-Niculescu-Mizil, C.; Sudhof, M.; Jurafsky, D.; Leskovec, J.; and Potts, C. 2013. A computational approach to politeness with application to social factors. *arXiv preprint arXiv:1306.6078*.

Derakhshan, A.; Mikaeili, M.; and Gedeon, T. 2018. Discriminating Between Truthfulness and Deception Using Infrared Thermal Imaging and Peripheral Physiology.

Ding, H.; Jiang, T.; and Riloff, E. 2018. Why is an Event Affective? Classifying Affective Events based on Human Needs. In *Proceedings of the AAAI-18 Workshop on Affective Content Analysis*. New Orleans, USA: AAAI.

Dragut, E. C.; Yu, C.; Sistla, P.; and Meng, W. 2010. Construction of a sentimental word dictionary. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, 1761–1764. New York, NY, USA: ACM.

Dumpala, S. H.; Chakraborty, R.; and Kopparapu, S. K. 2018. Knowledge driven feed-forward neural network for audio affective content analysis. In *Proceedings of the AAAI-18 Workshop on Affective Content Analysis*. New Orleans, USA: AAAI.

Ekman, P. 1992. An argument for basic emotions. *Cognition & Emotion* 6(3-4):169–200.

Esuli, A., and Sebastiani, F. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of LREC*, volume 6, 417–422. Citeseer.

Hatzivassiloglou, V., and McKeown, K. R. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, 174–181. Association for Computational Linguistics.

He, S.; Ballard, D.; and Gildea, D. 2004. Building and tagging with an affect lexicon.

Hu, C.; Walker, M. A.; Neff, M.; and Tree, J. E. F. 2015. Storytelling agents with personality and adaptivity. In *International Conference on Intelligent Virtual Agents*, 181–193. Springer.

Huang, C.-r. 2010. *Ontology and the Lexicon: A Natural Language Processing Perspective*. Cambridge University Press.

Jones, E.; Oliphant, T.; Peterson, P.; et al. 2001–. SciPy: Open source scientific tools for Python. [Online; accessed 2017-01-17].

Kanayama, H., and Nasukawa, T. 2006a. Fully automatic lexicon expansion for domain-oriented sentiment analysis. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, 355–363. Association for Computational Linguistics.

Kanayama, H., and Nasukawa, T. 2006b. Fully automatic lexicon expansion for domain-oriented sentiment analysis. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, EMNLP '06, 355–363. Stroudsburg, PA, USA: Association for Computational Linguistics.

Le, H.; Cerisara, C.; and Denis, A. 2018. Do convolutional networks need to be deep for Text Classification. In *Proceedings of the AAAI-18 Workshop on Affective Content Analysis*. New Orleans, USA: AAAI.

Litvinova, T.; Litvinova, O.; Seredin, P.; and Zagorovskaya, O. 2018. RusNeuroPsych: Corpus for Study Relations Between Author Demo-graphic, Personality Traits, Lateral Preferences and Affect in Text.

Liu, T., and Kappas, A. 2018. Predicting Engagement Breakdown in HRI Using Thin-slices of Facial Expressions. In *Proceedings of the AAAI-18 Workshop on Affective Content Analysis*. New Orleans, USA: AAAI.

Lu, Y.; Castellanos, M.; Dayal, U.; and Zhai, C. 2011. Automatic construction of a context-aware sentiment lexicon: An optimization approach. In *Proceedings of the 20th International Conference on World wide web*, 347–356. ACM.

Maheshwari, T.; Reganti, A. N.; Kumar, U.; Chakraborty, T.; and Das, A. 2017. Semantic interpretation of social network communities. In *AAAI*, 4967–4968.

Mairesse, F., and Walker, M. A. 2008. Trainable generation of big-five personality styles through data-driven parameter estimation. In *ACL*, 165–173.

Manning, C. D.; Surdeanu, M.; Bauer, J.; Finkel, J. R.; Bethard, S.; and McClosky, D. 2014. The stanford corenlp natural language processing toolkit. In *ACL (System Demonstrations)*, 55–60.

Mehrabian, A. 1980. Basic dimensions for a general psychological theory implications for personality, social, environmental, and developmental studies.

Miller, G. A. 1995. Wordnet: a lexical database for english. *Communications of the ACM* 38(11):39–41.

Mitchell, S.; OSullivan, M.; and Dunning, I. 2011. Pulp: A linear programming toolkit for python. *The University of Auckland, Auckland, New Zealand, http://www.optimization-online. org/DB_FILE/2011/09/3178. pdf*.

Mohammad, S. M., and Turney, P. D. 2010. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 workshop on computational approaches to analysis and generation of emotion in text*, 26–34. Association for Computational Linguistics.

Mohammad, S. M., and Turney, P. D. 2013a. Crowdsourcing a word-emotion association lexicon. 29(3):436–465.

Mohammad, S. M., and Turney, P. D. 2013b. Nrc emotion lexicon. Technical report, NRC Technical Report.

Mohammad, S. M. 2012. # emotional tweets. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*, 246–255. Association for Computational Linguistics.

Muhammad, A.; Wiratunga, N.; Lothian, R.; and Glassey, R. 2013. Domain-based lexicon enhancement for sentiment analysis. In *SMA@ BCS-SGAI*, 7–18.

Nester, D.; Haduong, N.; Vaidyanathan, P.; Prud'Hommeaux, E.; Bailey, R.; and Alm, C. 2018. Multimodal Alignment for Affective Content. In *Proceedings of the AAAI-18 Workshop on Affective Content Analysis*. New Orleans, USA: AAAI.

Nielsen, F. Å. 2011. A new anew: Evaluation of a word list for sentiment analysis in microblogs. *arXiv preprint arXiv:1103.2903*.

Pavlick, E., and Tetreault, J. 2016. An empirical analysis of formality in online communication. *Transactions of the Association for Computational Linguistics* 4:61–74.

Pennebaker, J. W.; Francis, M. E.; and Booth, R. J. 2001. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates* 71:2001.

Peterson, K.; Hohensee, M.; and Xia, F. 2011. Email formality in the workplace: A case study on the enron corpus. In *Proceedings of the Workshop on Languages in Social Media*, 86–95. Association for Computational Linguistics.

Picard, R. W., and Picard, R. 1997. *Affective computing*, volume 252. MIT press Cambridge.

Pittman, M., and Reich, B. 2016. Social media and loneliness: Why an instagram picture may be worth more than a thousand twitter words. *Computers in Human Behavior* 62:155–167.

Plutchik, R. 1980. A general psychoevolutionary theory of emotion. *Theories of Emotion* 1:3–31.

Preotiuc-Pietro, D.; Xu, W.; and Ungar, L. H. 2016. Discovering user attribute stylistic differences via paraphrasing. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, 3030–3037.

Qiu, G.; Liu, B.; Bu, J.; and Chen, C. 2011. Opinion word expansion and target extraction through double propagation. *Computational linguistics* 37(1):9–27.

Rahimtoroghi, E.; Wu, J.; Wang, R.; Anand, P.; and Walker, M. A. 2017. Modelling protagonist goals and desires in first-person narrative. *arXiv preprint arXiv:1708.09040*.

Ranganath, R.; Jurafsky, D.; and McFarland, D. A. 2013. Detecting friendly, flirtatious, awkward, and assertive speech in speed-dates. *Computer Speech & Language* 27(1):89–115.

Reed, L.; Wu, J.; Oraby, S.; Anand, P.; and Walker, M. 2017. Learning lexico-functional patterns for first-person affect. *arXiv preprint arXiv:1708.09789*.

Ribeiro, F. N.; Araújo, M.; Gonçalves, P.; Gonçalves, M. A.; and Benevenuto, F. 2016. Sentibench-a benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Science* 5(1):1–29.

Robertson, S. 2004. Understanding inverse document frequency: On theoretical arguments for idf. *Journal of Documentation* 60(5):503–520.

Sakaguchi, K.; Duh, K.; Post, M.; and Van Durme, B. 2017. Robsut wrod reocginiton via semi-character recurrent neural network. In *AAAI*, 3281–3287.

Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; Seligman, M. E.; et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS One* 8(9):e73791.

Seligman, M. E. 2011. Flourish: a visionary new understanding of happiness and well-being. *Policy* 27(3):60–1.

Stackoverflow. 2015. Finding the best trade-off point on a curve.

Stone, P. J.; Dunphy, D. C.; and Smith, M. S. 1966. The general inquirer: A computer approach to content analysis.

Strapparava, C., and Mihalcea, R. 2007. Semeval-2007 task 14: Affective text. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, 70–74. Association for Computational Linguistics.

Strapparava, C.; Valitutti, A.; et al. 2004. Wordnet affect: an affective extension of wordnet. In *LREC*, volume 4, 1083–1086.

Taboada, M.; Brooke, J.; Tofiloski, M.; Voll, K.; and Stede, M. 2011. Lexicon-based methods for sentiment analysis. *Computational Linguistics* 37(2):267–307.

Thelwall, M.; Buckley, K.; Paltoglou, G.; Skowron, M.; Garcia, D.; Gobron, S.; Ahn, J.; Kappas, A.; Küster, D.; and Holyst, J. A. 2013. Damping sentiment analysis in online communication: discussions, monologs and dialogs. In *International Conference on Intelligent Text Processing and Computational Linguistics*, 1–12. Springer.

Tighe, E. P.; Ureta, J. C.; Pollo, B. A. L.; Cheng, C. K.; and de Dios Bulos, R. 2016. Personality trait classification of essays with the application of feature reduction. In *SAAIP@ IJCAI*, 22–28.

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

Warriner, A. B.; Kuperman, V.; and Brysbaert, M. 2013. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods* 45(4):1191–1207.

Wiebe, J.; Wilson, T.; and Bell, M. 2001. Identifying collocations for recognizing opinions. In *Proceedings of the ACL-01 Workshop on Collocation: Computational Extraction, Analysis, and Exploitation*, 24–31.

Wu, J.; Walker, M.; Anand, P.; and Whittaker, S. 2017. Linguistic reflexes of well-being and happiness in echo. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 81–91.

Zahiri, S., and Choi, J. 2018. Emotion Detection on TV Show Transcripts with Sequence-based Convolutional Neural Networks. In *Proceedings of the AAAI-18 Workshop on Affective Content Analysis*. New Orleans, USA: AAAI.

Zhao, S.; Yao, H.; and Jiang, X. 2015. Predicting continuous probability distribution of image emotions in valence-arousal space. In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM '15, 879–882. New York, NY, USA: ACM.