DR. SAIF M. MOHAMMAD

Last updated: August 02, 2023

Website: http://www.saifmohammad.com
Google Scholar Profile: http://www.purl.org/net/googlescholar

Email: saif.mohammad@nrc-cnrc.gc.ca
Twitter: @SaifMMohammad

Address: 1200 Montreal Road, Building M-50,
Ottawa, Ontario, K1A 0R6, Canada.

AREAS OF INTEREST

Computational Linguistics, Natural language processing (NLP), Lexical Semantics, Emotions in Language, Sentiment Analysis, AI Ethics, Computational Creativity, Information Visualization, NLP for Psychology, Computational Social Science.

CURRENT POSITION

Senior Research Scientist, October 2014–present
Information and Communications Technologies, National Research Council Canada, Ottawa, Canada

PREVIOUS POSITIONS

Research Scientist, October 2009–September 2014
Institute of Information Technology, National Research Council Canada, Ottawa, Canada

Research Associate, January 2008–September 2009
Institute of Advanced Computer Studies, University of Maryland, College Park, USA

EDUCATION

Doctor of Philosophy, September 2003–February 2008
Department of Computer Science, University of Toronto, Toronto, Canada

Master of Science, September 2001–August 2003
Department of Computer Science, University of Minnesota, Duluth, USA

Bachelor of Engineering, August 1996–May 2000
Department of Electronics Engineering, University of Pune, Pune, India

HIGHLIGHTS

- #Articles published 117; h-index: 57; #citations: 20,198 ($\sim$3K citations in 2022 alone)
- Best Paper Award – Honorable Mention. *Examining the Citational Amnesia in NLP*, ACL 2023.
- Best Paper Award. *AfriSenti: A Benchmark Twitter Sentiment Analysis for African Languages*, 4th AfricaNLP Workshop, ICLR 2023.
- Research Excellence Award, National Research Council Canada, 2022.
- Chair of the 2019 Canada–UK Symposium on Ethics in AI, Feb 21–22, Ottawa, Canada.
- Action Editor, Computational Linguistics: 2020–2023; Associate Editor, Journal of Artificial Intelligence Research: 2023–2024; Action Editor, Inaugural Senior Action Editor, ACL Rolling Review 2022–2024; Action Editor, Transactions of ACL 2022–2024
- Senior Area Chair (SAC) for EMNLP 2023, Area Chair (AC) for EACL 2023, ACL 2022, EMNLP 2022, SAC for ACL 2021, AC for EMNLP 2020, ACL 2020, *Sem 2020, ACL 2019, EMNLP 2019, *Sem 2019, ACL 2017, EMNLP 2017, EMNLP 2016, and NAACL 2015 in the broad areas of Sentiment Analysis, Opinion Mining, Argument Mining, Ethics, and Lexical Semantics. AC for the first Ethics, Bias and Fairness track at NAACL-2019.

- Workshops Chair for ACL 2020; Co-chair of SemEval-2019, SemEval 2018, and SemEval-2017. (SemEval is one of the largest NLP evaluation exercises with proceedings as large as that of top NLP conferences, and about 1000 participants each year.)
- Co-organizer of the 8th and 9th Workshops on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA-2017 and WASSA-2018).
- Organized international competitions on previously unexplored tasks such as the SemEval-2023 Task 12 (first NLP shared task on African languages), SemEval-2018 Task-1 (the first NLP shared task with a separate evaluation component on fairness and bias), and the SemEval-2016 Task-6 (First shared task on Detecting Stance from Tweets).
- Led the NRC-Canada team in developing a sentiment analysis system that ranked first in recent international shared tasks (among more than 40 submissions).
- Keynote speeches at:
    - 2022 Social NLP Workshop at ACL 2022, July 14, 2022, Seattle, USA.
    - 2019 Canada–UK Workshop on Practical AI Ethics, March 12, 2019, London, UK.
    - 6th evaluation campaign EVALITA 2018. December 12, 2018. Turin, Italy.
    - ACL 2018 Workshop on NLP for Social Media, Melbourne, Australia.
    - International Conference on Asian Language Processing (IALP 2016), Tainan, Taiwan.
    - the COLING 2016 workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media (PEOPLES 2016), Osaka, Japan.
    - the 2016 IR Workshop at the Academia Sinica, November 24, 2016, Taipei, Taiwan
    - ACL 2014 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)
- Featured speaker at the Distinguished Computational Linguistics Lecture at Rochester Institute of Technology. March 1, 2018. Rochester, NY.
- Presented tutorials at international NLP conferences on Sentiment and Emotion Analysis. A video of the EMNLP-2014 tutorial has received more than 23,000 views on YouTube. (https://www.youtube.com/watch?v=zv16Xyph7Ss)
- Work cited in several articles in the popular press, including articles in Time, Slashdot, LiveScience, io9, The Physics arXiv Blog, PC World, and Popular Science.
- Helped develop TransProse, a system for generating music that captures the emotions in a piece of literature. The system uses the NRC Emotion Lexicon.
    - A symphony orchestra performed music composed using TransProse and a human composer under the glass of the Louvre museum in Paris, September 20, 2016.
- An art project, the Wishing Wall, that uses the NRC Emotion lexicon was displayed in:
    - Tekniska Museet, Stockholm, Sweden (Oct 14 – Aug 15)
    - Onassis Cultural Centre, Athens (19th Oct'15 – 10th Jan'16)
- Recognition: *Top Writer for Science on Medium*. November 2019. (One of 50 at the time. There were about 63K stories in the Science category on Medium at the time.)

PROFESSIONAL ACTIVITIES

Positions Held at International Conferences and Journals:

- 2023
  - Action Editor, Computational Linguistics (CL).
  - Action Editor, Transactions in ACL (TACL).
  - Associate Editor, Journal of Artificial Intelligence Research (JAIR).
  - Senior Action Editor, ACL Rolling Review.
  - Senior Area Chair for Sentiment Analyisis and Argument Mining, EMNLP 2023.
  - Area Chair for Ethics and Sustainable NLP, EACL 2023.
- 2022
  - Action Editor, Computational Linguistics (CL).
  - Action Editor, Transactions in ACL (TACL).
  - Senior Action Editor, ACL Rolling Review.
  - Area Chair for Ethics and NLP, ACL 2022.
  - Area Chair for Open Questions, Major Obstacles, and Unresolved Issues in NLP, EMNLP 2022.
- 2021
  - Action Editor, Computational Linguistics.
  - Action Editor, ACL Rolling Review.
  - Senior Area Chair for Sentiment Analysis, Stylistic Analysis, and Argument Mining, ACL 2021.
- 2020
  - Action Editor, Computational Linguistics.
  - Senior Area Chair for Sentiment Analysis, Stylistic Analysis, and Argument Mining, EMNLP 2020.
  - Area Chair for Sentiment Analysis, Stylistic Analysis, and Argument Mining, StarSem 2020, Barcelona, Spain.
  - Area chair for Lexical Semantics, ACL 2020.
  - Workshops Chair ACL 2020.
- 2019
  - Chair of the 2019 Canada–UK Symposium on Ethics in AI, Ottawa, Canada.
  - Co-chair of SemEval 2019, Minneapolis, Minnesota, USA.
  - Area chair for Sentiment Analysis and Argumentation Mining, EMNLP-IJCNLP 2019, Hong Kong.

- Area chair for Sentiment Analysis and Argumentation Mining, ACL 2019, Florence, Italy.
- Area chair for Ethics, Bias and Fairness, NAACL 2019, Minneapolis, Minnesota, USA.
- Area chair for Semantics in NLP Applications, *Sem 2019, Minneapolis, Minnesota, USA. Mentor, Student Research Workshop, ACL 2019, Florence, Italy.
- Mentor, Student Research Workshop, NAACL 2019, Minneapolis, Minnesota, USA.

- 2018
  - Co-chair of SemEval 2018, New Orleans, Louisiana, USA.
  - Co-organizer of WASSA 2018, Brussels, Belgium.
  - Mentor, Student Research Workshop, NAACL 2018, New Orleans, Louisiana, USA.
  - Mentor, Student Research Workshop, ACL 2018, Melbourne, Australia.

- 2017
  - Co-chair of SemEval 2017, Vancouver, Canada.
  - Co-organizer of WASSA 2018, Copenhagen, Denmark.
  - Area chair for Sentiment and Opinion Mining, ACL 2017, Vancouver, Canada.
  - Area chair for Sentiment and Opinion Mining, EMNLP 2017, Copenhagen, Denmark.
  - Mentor, Student Research Workshop, ACL 2017, Vancouver, Canada.

- 2016
  - Area chair for Sentiment and Opinion Mining, EMNLP 2016, Austin, Texas, USA.
  - Publicity chair, EMNLP 2016, Austin, Texas, USA.

- 2015
  - Area chair for Sentiment and Opinion Mining, NAACL 2015, Denver, Colorado.
  - Publicity chair, NAACL 2015, Denver, Colorado.
  - Mentor, Student Research Workshop, NAACL 2015, Denver, Colorado.

- 2014, 2012, 2011
  - Mentor, Student Research Workshops, ACL 2014, NAACL 2012, ACL 2011.

- Most years 2008–present
  - Program committee member for top NLP journals and conferences, including TACL, ACL, NAACL, EMNLP, CL, CI, NLE, JAIR, TSLP, IJCNLP, ICWSM, and IJCAI.

Tutorials Presented at International Conferences:

- Computational Analysis of Affect and Emotion in Language. Saif M. Mohammad and Cecilia Ovesdotter Alm. Tutorial at the 2015 Conference on Empirical Methods on Natural Language Processing, September 2015, Lisboa, Portugal.

- Sentiment Analysis of Social Media Texts. Saif M. Mohammad and Xiaodan Zhu. Tutorial at the 2014 Conference on Empirical Methods on Natural Language Processing, October 2014, Doha, Qatar.

Organization of Shared Task Competitions

1. Co-organizer of SemEval-2023 Task 12: Sentiment Analysis for African Languages.
2. Co-organizer of SemEval-2018 Task 1: Affect in Tweets. 75 teams and about 200 participants.
3. Co-organizer of WASSA-2018 Implicit Emotions Shared Task.
4. Co-organizer of shared task on detecting emotion intensity at WASSA-2017.
5. Co-organizer of Semeval-2016 Task 6: Detecting Stance in Tweets.
6. Co-organizer of Semeval-2016 Task 7: Determining Sentiment Intensity of English and Arabic Phrases.
7. Co-organizer of SemEval-2015 Task 10: Sentiment Analysis in Twitter.
8. Co-organizer of SemEval-2012 Task 2: Measuring Degrees of Relational Similarity.

Submissions to Shared Task Competitions

1. AMIA Shared Task on detecting adverse drug reactions in tweets, Washington, DC, USA, 2017.
   Task: Classification of tweets mentioning adverse drug reactions.
   Team: Svetlana Kiritchenko, Saif M. Mohammad, Jason Morin, and Berry de Bruijn.
   Result: Placed first. (9 teams participated.)
2. SemEval-2014 Task 4: Aspect Based Sentiment Analysis, August 2014, Dublin, Ireland.
   Task: Determine sentiment towards aspect terms and aspect categories.
   Team: NRC-Canada: Saif Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu.
   Result: Placed first in two of the three sentiment tasks. (30+ teams participated.)
3. SemEval-2014 Task 9: Sentiment Analysis in Twitter, August 2014, Dublin, Ireland.
   Task: Determine sentiment of tweets, SMS messages, and blog posts.
   Team: NRC-Canada: Saif Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu.
   Result: Placed first in five of ten sub-tasks. (40+ teams participated.)
4. SemEval-2013 Task 2: Sentiment Analysis in Twitter, June 2013, Atlanta, USA.
   Task: Determine sentiment of tweets and SMS messages.
   Team: NRC-Canada: Saif Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu.
   Result: Placed first in three of four sub-tasks. (40+ teams participated.)
5. Automatic Content Extraction Evaluation (ACE-08), September 2008, Baltimore, MD.
   Task: Multidocument coreference resolution.
   Team: Johns Hopkins HLT/COE team.
   Result: Placed third (behind IBM and BBN Technologies).

Information Visualization and Data Sonification Demos

1. Citational Amnesia: An Online Demo that Shows Information about the Age Diversity of Citations of a Paper
   https://huggingface.co/spaces/ mrungta8/CitationalAmnesia/

2. *NLP Scholar*: An Interactive Visualization to Explore NLP Papers
   http://saifmohammad.com/WebPages/nlpscholar-demo-basic.html

   A visual explorer to help users find related work and explore broad patterns in citations.

3. *TransProse: Converting Text to Music.* Hannah Davis and Saif M. Mohammad.
   http://www.musicfromtext.com

   A system that takes as input classic English novels and generates music that captures the flow of emotions in it.

4. *ImagiSaurus: An Interactive Visualization of the Roget's Thesaurus*
   http://www.saifmohammad.com/WebPages/imagisaurus.html

   Imagisaurus helps users quickly grasp the nature and size of the thesaurus taxonomy. Additionally, Imagisaurus allows exploration of affectual categories in the thesaurus.

5. *Lexichrome: An Interactive Word–Color Catalogue for Scholars, Designers, and Writers.* Chris Kim, Saif M. Mohammad, and Christopher Colins.
   http://www.lexichrome.com/#palette

   Lexichrome is an application that maps the NRC Word–Colour Association Lexicon to a web-based catalogue that users can browse and analyze. It allows visitors to view the chromatic makeup (color associations) of a user-provided text.

6. *An Interactive Visualizer for the NRC Emotion Lexicon*
   http://www.saifmohammad.com/WebPages/NRC-Emotion-Lexicon.htm

   The NRC Emotion Lexicon is a list of English words and their associations with eight basic emotions (anger, fear, anticipation, trust, surprise, sadness, joy, and disgust) and two sentiments (negative and positive). The annotations were manually done by crowdsourcing.

7. *An Interactive Visualizer for Sentiment Composition Lexicons*
   http://www.saifmohammad.com/WebPages/SCL.html

   Sentiment composition is the determining of sentiment of a multi-word linguistic unit, such as a phrase or a sentence, based on its constituents. We present two visualizations: (1) exploring sentiment composition in phrases formed by at least one positive and at least one negative word—phrases like happy accident and best winter break, and (2) exploring sentiment composition in phrases formed with negators, modals, and degree adverbs.

8. *An Interactive Visualizer for the Stance Dataset*
   http://www.saifmohammad.com/WebPages/StanceDataset.htm

   We visualize a dataset of tweets manually annotated for stance towards given target, target of opinion (opinion towards), and sentiment (polarity).

INVITED TALKS

1. The Search for Emotions, Creativity, and Fairness in Language. Ludwig Maximilian University of Munich, May 8, 2023.

2. The Search for Emotions, Creativity, and Fairness in Language. University of Gottingen, Dec 14, 2022.

3. The Search for Emotions, Creativity, and Fairness in Language. Keynote talk at the Emotion Workshop, University of Ghent, Dec 1, 2022.

4. The Search for Emotions, Creativity, and Fairness in Language. The Statistical Society of Canada (SSC), Webinar, Sep 4, 2022.

5. Ethics Sheets for Social NLP. The 2022 Social NLP Workshop at ACL 2022, July 14, 2022, Seattle, USA.

6. The Search for Emotions, Creativity, and Fairness in Language. Defence Research and Development Canada, Toronto Research Centre, April 21, 2022, Toronto, Canada.

7. The Search for Emotions in Language. Carolina Affective Science Lab. University of North Carolina. Feb 21, 2022.

8. Emotions in Text: The Power and Versatility of Large Emotion Lexicons. Canadian Digital Services. January 13, 2022.

9. Judging AI: How to Develop a Healthy Mistrust of AI. The Canadian Appellate Courts Seminar, Halifax, Canada. Nov 25, 2021.

10. Ethics Sheets for AI Tasks and a Case Study for Automatic Emotion Recognition. The University of British Columbia Language Sciences Talks, Vancouver, Canada. July 15, 2021.

11. Emotion Dynamics of Fictional Characters using Large Emotion Lexicons. The 17th IGEL Conference. June 23, 2021.

12. Ethics Sheets for AI Tasks and a Case Study for Automatic Emotion Recognition. The Alan Turing Institute, London, UK. May, 2021.

13. Gender Gap in Natural Language Processing Research: Disparities in Authorship and Citations. Department of Computer Science, University of Sheffield. October 2020. Sheffield, UK.

14. The Search for Emotions, Creativity, and Fairness in Language. Data Systems Seminar Series, University of Waterloo, October 28 2019, Waterloo, Canada.

15. Fairness and Emotions in Language. The Globe and Mail. October 29, 2019, Toronto, Canada.

16. Creativity and Emotions in Language. Invited talk and panel on AI and Creativity in Government at the Creative Marketplace Lab on Data, Skills and Technology, Department of Canadian Heritage, September 30 2019, Gatineau, Canada.

17. The Search for Emotions, Creativity, and Fairness in Language. Mila - Quebec AI Institute, August 2019, Montreal, Canada.

18. The Search for Emotions, Creativity, and Fairness in Language. Keynote speech at the 8th KDD Workshop on Issues of Sentiment Discovery and Opinion Mining (WISDOM-2019), August 2019, Anchorage, Alaska.

19. Examining Fairness in Language Through Emotions. Keynote speech at the 2019 Canada–UK Workshop on Practical AI Ethics, March 12, 2019, London, UK.

20. The Search for Emotions, Creativity, and Fairness in Language. The Alan Turing Institute. March 11, 2019. London, UK.

21. Examining Fairness in Language Through Emotions. The 2019 Canada–UK Symposium on Ethics in AI, Feb 21, 2019, Ottawa, Canada.

22. The Search for Emotions, Creativity, and Fairness in Language. Keynote speech at the 6th evaluation campaign EVALITA 2018. December 12, 2018. Turin, Italy.

23. The Search for Emotions, Creativity, and Fairness in Language. Department of Computational Linguistics, Heidelberg University. December 13, 2018. Heidelberg, Germany.

24. The Search for Emotions, Creativity, and Fairness in Language. Invited Speaker Series. Pickering Centre for Human Development, Carleton University. November 16, 2018. Ottawa, Canada.

25. The Search for Emotions, Creativity, and Fairness in Language. Short talk to 9th graders on the occasion of Bring-Your-Ninth-Grader-to-Work Day at the National Research Council Canada. November 14, 2018. Ottawa, Canada.

26. The Search for Emotions, Creativity, and Fairness in Language. At the Natural Language Processing Seminar and AI Seminar, University of Michigan, Ann Arbor. October 2, 2018. Ann Arbor, Michigan.

27. The Search for Emotions, Creativity, and Fairness in Language. Department of Computer Science, University of Utah. August 6, 2018. Salt Lake City, Utah.

28. The Search for Emotions in Language. Keynote speech at the ACL-2018 Workshop on Natural Language Processing for Social Media, July 2018, Melbourne, Australia.

29. The Search for Emotions in Language. The University of Melbourne, July 2018, Melbourne, Australia.

30. The Search for Emotions in Language. The University of Waikato, July 2018, Hamilton, New Zealand.

31. Examining Gender and Race Bias in Two Hundred Sentiment Analysis Systems. Second ACL Workshop on Ethics in Natural Language Processing, New Orleans, LA, USA, June 2018.

32. The Search for Emotions in Language. Google Research. May 25, 2018. Montreal, Canada.

33. The Search for Emotions in Language. Featured speaker at the Distinguished Computational Linguistics Lecture at Rochester Institute of Technology. March 1, 2018. Rochester, NY.

34. The Search for Emotions in Language. University of Toronto. Feb 27, 2018. Toronto, Canada.

35. Finding Emotions in Language. Information Systems Sprott School of Business, Carleton University. November 13, 2017. Ottawa, Canada.
36. Emotion-Aware Machines. Short talk to 9th graders on the occasion of Bring-Your-Ninth-Grader-to-Work Day at the National Research Council Canada. November 1, 2017. Ottawa, Canada.
37. Affect Associations in Creative Language. Seminar Series for the Doctorate in Cognitive Informatics at Université du Qubéc à Montrál (UQAM), April 13, 2017, Montreal, Canada.
38. Affect Associations in Creative Language. Keynote speech at the COLING 2016 workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media (PEOPLES 2016), December 12, 2016, Osaka, Japan.
39. Affect Associations: The Building Blocks of Sentiment Analysis. Keynote speech at the 20th International Conference on Asian Language Processing (IALP 2016), November 21-23, 2016, National Cheng Kung University, Tainan, Taiwan.
40. Affect Associations in Creative Language. Keynote speech at the 2016 IR Workshop at the Academia Sinica, November 24, 2016, Taipei, Taiwan.
41. Capturing Reliable Fine-Grained Sentiment Associations by Crowdsourcing and Best–Worst Scaling. Invited talk at the TAMALE seminar series, September 15, 2016, University of Ottawa. (Joint work with, and talk given by, Svetlana Kiritchenko.)
42. Metaphor as a Medium for Emotion: An Empirical Study. Invited talk at the TAMALE seminar series, September 15, 2016, University of Ottawa. (Joint work with Ekaterina Shutova and Peter Turney.)
43. Sentiment and Emotions in Tweets. Department of Geography and Geoinformation Science at George Mason University. May 29, 2016, FairFax, Virginia.
44. The Words are Alive: Associations with Sentiment, Emotions, Colours, and Music. Language Technology Institute, Carnegie Mellon University, September 2014, Pittsburgh, PA.
45. Words: Evaluative, Emotional, Colourful, Musical! Keynote speech at the ACL 2014 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA), June 2014, Baltimore, MD.
46. Crowdsourcing Word-Emotion Associations. October 2013, University of Pennsylvania, Philadelphia, PA.
47. Sentiment Analysis with Social Media Text. June 2013, MediaMiser, Ottawa, Canada.
48. Sentiment Analysis with Social Media Text. June 2013, Thales, Ottawa, Canada.
49. Analyzing Electoral Tweets for Affect, Style, and Purpose. April 2013, University of Ottawa.
50. From Once Upon a Time to Happily Ever After: Tracking Emotions in Books and Mail. July 2011: Digital Books and Social Media groups, Amazon Research, Seattle, OR, USA.
51. From Once Upon a Time to Happily Ever After: Tracking Emotions in Books and Mail. June 2011: Defence Research and Development Canada (DRDC) workshop, Ottawa, Canada.
52. Crowdsourcing the creation of an emotion lexicon. August 2010: Microsoft Research, Bangalore, India.
53. Crowdsourcing the creation of an emotion lexicon. August 2010: IBM Research, Bangalore, India.

54. Crowdsourcing the creation of an emotion lexicon. July 2010: Department of Computational Linguistics at the University of Heidelberg and the NLP group from HITS gGmbH, Heidelberg, Germany.
55. Crowdsourcing the creation of an emotion lexicon. July 2010: Department of Computer Science, Technische Universität Darmstadt, Darmstadt, Germany.
56. Hybrid approaches to semantic closeness and oppositeness. June 2009: The Text Analysis and Machine Learning Group, University of Ottawa, Ottawa, Canada.
57. Hybrid approaches to semantic closeness and oppositeness. June 2009: National Research Council - The Institute for Information Technology (NRC-IIT), Ottawa, Canada.
58. Hybrid approaches to semantic closeness and semantic oppositeness. October 2008: Text Mining Sciences, Yahoo! Research, Santa Clara, CA.
59. Hybrid approaches to semantic closeness and semantic oppositeness. October 2008: Human Language Technology Research Institute, University of Texas at Dallas, Dallas, TX.
60. Hybrid approaches to semantic closeness and semantic oppositeness. October 2008: University of North Texas, Dallas, TX.
61. Determining Degree of Antonymy: Step 1 for Detecting Erroneous Coreference Links by Contradiction. March 2008. *Human Language Technology Center of Excellence (HLT/COE) Quarterly Technical Exchange*, Johns Hopkins HLT/COE, Baltimore, MD.

PUBLICATIONS

Summary

h-index: 57
Number of publications: 117
Number of citations received: 20,198 ($\sim$3K citations in 2022 alone)
Google Scholar Profile: http://www.purl.org/net/googlescholar

Book Chapters

1. Saif M. Mohammad. 2020. Sentiment Analysis: Automatically Detecting Valence, Emotions, and Other Affectual States from Text. Book chapter in *The 2nd Edition of Emotion Measurement*, Elsevier.

2. Saif M. Mohammad. 2016. Challenges in Sentiment Analysis. Book chapter in *A Practical Guide to Sentiment Analysis*, Springer.

3. Saif M. Mohammad. 2016. Sentiment Analysis: Detecting Valence, Emotions, and Other Affectual States from Text. Book chapter in *Emotion Measurement*, Elsevier.

4. Graeme Hirst and Saif M. Mohammad. 2009. Semantic distance measures with distributional profiles of coarse-grained concepts. Book chapter in *Modeling, Learning and Processing of Text Technological Data Structures*, Editors Alexander Mehler, Kai-Uwe Kühnberger, Henning Lobin, Harald Lüngen, Angelika Storrer, and Andreas Witt. Springer.

Refereed Journal Articles

1. Saif M. Mohammad. Ethics Sheet for Automatic Emotion Recognition and Sentiment Analysis. Computational Linguistics, 48 (2): 239–278. June 2022.

2. Will E. Hipson, Svetlana Kiritchenko, Robert J. Coplan, Saif M. Mohammad. Examining the Language of Solitude vs. Loneliness in Tweets. Journal of Social and Personal Relationships. March 2021.

3. Felipe Bravo-Marquez, Eibe Frank, Bernhard Pfahringer, Saif M. Mohammad. 2019. AffectiveTweets: a Weka Package for Analyzing Affect in Tweets. Journal of Machine Learning Research, 20(92):1–6.

4. Peter D. Turney and Saif M. Mohammad. 2019. The Natural Selection of Words: Finding the Features of Fitness. *PLoS One*, 14 (1):e0211512.

5. Abeed Sarker, Maksim Belusov, Jasper Friedrichs, Kai Hakala, Sifei Han, Svetlana Kiritchenko, Farrokh Mehryary, Anthony Rios, Tung Tran, Berry de Bruijn, Filip Ginter, Ramakanth Kavuluru, Debanjan Mahata, Saif M. Mohammad, Goran Nenadic, Graciela Gonzalez-Hernandez. 2018. Data and systems for medication-related text classification and concept normalization from Twitter: Insights from the Social Media Mining for Health (SMM4H)-2017 shared task. Journal of the American Medical Informatics Association (JAMIA). 25(10), 1274–1283.

6. Saif M. Mohammad, Parinaz Sobhani, and Svetlana Kiritchenko. 2017. Stance and Sentiment in Tweets. *Special Section of the ACM Transactions on Internet Technology on Argumentation in Social Media*, 17(3).

7. Preslav Nakov, Sara Rosenthal, Svetlana Kiritchenko, Saif M. Mohammad, Zornitsa Kozareva, Alan Ritter, Veselin Stoyanov, and Xiaodan Zhu. 2016. Developing a Successful SemEval Task in Sentiment Analysis of Twitter and Other Social Media Texts. *Language Resources and Evaluation*. 50(1):1–31.

8. Saif M. Mohammad, Mohammad Salameh, and Svetlana Kiritchenko. 2016. How Translation Alters Sentiment. *Journal of Artificial Intelligence Research*, Volume 55, pages 95–130.

9. Peter D. Turney and Saif M. Mohammad. 2015. Experiments with Three Approaches to Recognizing Lexical Entailment. *Natural Language Engineering*, 21(3):437–476.

10. Saif M. Mohammad, Svetlana Kiritchenko, Xiaodan Zhu, and Joel Martin. 2015. Sentiment, Emotion, Purpose, and Style in Electoral Tweets. *Information Processing and Management*, Volume 51, Issue 4, Pages 480–499.

11. Saif M. Mohammad and Svetlana Kiritchenko. 2015. Using Hashtags to Capture Fine Emotion Categories from Tweets. *Computational Intelligence*, 31(2):301–326.

12. Svetlana Kiritchenko, Xiaodan Zhu and Saif M. Mohammad. 2014. Sentiment Analysis of Short Informal Texts. *Journal of Artificial Intelligence Research*, volume 50, pages 723–762.

13. Saif M. Mohammad, Bonnie Dorr, Graeme Hirst, and Peter Turney. 2013. Computing Lexical Contrast. *Computational Linguistics*, 39(3), 555–590.

14. Saif M. Mohammad and Peter Turney. 2013. Crowdsourcing a Word–Emotion Association Lexicon. *Computational Intelligence*, 29(3), 436–465.

15. Vahed Qazvinian, Dragomir R. Radev, Saif M. Mohammad, Bonnie Dorr, David Zajic, Michael Whidby, and Taesun Moon. 2013. Generating Extractive Summaries of Scientific Paradigms. *Journal of Artificial Intelligence Research (JAIR)*, 46, pages 165–201.

16. Colin Cherry, Saif M. Mohammad, and Berry de Bruijn. 2012. Binary Classifiers and Latent Sequence Models for Emotion Detection in Suicide Notes. *Journal of Biomedical Informatics Insights*, 5:147–154.

17. Saif M. Mohammad. 2012. From Once Upon a Time to Happily Ever After: Tracking Emotions in Mail and Books. Decision Support Systems, 53(4):730–741.

Refereed Conference Papers

1. Janvijay Singh, Mukund Rungta, Diyi Yang, and Saif M. Mohammad. Forgotten Knowledge: Examining the Citational Amnesia in NLP. In Proceedings of the 61st Annual Meeting of the Association of Computational Linguistics (ACL-2023), Toronto, Canada.

2. Mohamed Abdalla, Jan Philip Wahle, Terry Ruas, Aurélie Névéol, Fanny Ducel, Saif M. Mohammad, and Karen Fort. The Elephant in the Room: Analyzing the Presence of Big Tech in Natural Language Processing Research. In Proceedings of the 61st Annual Meeting of the Association of Computational Linguistics (ACL-2023), Toronto, Canada.

3. Jan Philip Wahle, Terry Ruas, Saif M. Mohammad, Norman Meuschke, and Bela Gipp. AI Usage Cards: Responsibly Reporting AI-generated Content. In Proceedings of JCDL 2023, Santa Fe, NM, USA.

4. Mohamed Abdalla, Krishnapriya Vishnubhotla, and Saif M. Mohammad. What Makes Sentences Semantically Related? A Textual Relatedness Dataset and Empirical Study. In Proceedings of EACL 2023, Dubrovnik, Croatia.

5. Saif M. Mohammad. Best Practices in the Creation and Use of Emotion Lexicons. In Proceedings of EACL 2023, Dubrovnik, Croatia.

6. Mukund Rungta, Janvijay Singh, Saif M. Mohammad, Diyi Yang. Geographic Citation Gaps in NLP Research. EMNLP, 2022, Abu Dhabi, UAE.

7. Adam Hammond, Krishnapriya Vishnubhotla, Graeme Hirst, and Saif M. Mohammad. Voices Speaking To and About One Another: Introducing the Project Dialogism Novel Corpus. In Proceedings of the Digital Humanities 2022 Conference, July 2022, virtual.

8. Krishnapriya Vishnubhotla and Saif M. Mohammad. Tweet Emotion Dynamics: Emotion Word Usage in Tweets from US and Canada. In Proceedings of the 13th Language Resources and Evaluation Conference (LREC-2022), May 2022, Marseille, France.

9. Jan Philip Wahle, Terry Ruas, Saif Mohammad and Bela Gipp. D3: A Massive Dataset of Scholarly Metadata for Analyzing the State of Computer Science Research. In Proceedings of the 13th Language Resources and Evaluation Conference (LREC-2022), May 2022, Marseille, France.

10. Saif M. Mohammad. Ethics Sheets for AI Tasks. In Proceedings of the 60th Annual Meeting of the Association of Computational Linguistics (ACL-2022), May 2022, Dublin, Ireland.

11. Rishav Hada, Sohi Sudhir, Pushkar Mishra, Helen Yannakoudakis, Saif M. Mohammad, and Ekaterina Shutova. Ruddit: Norms of Offensiveness for English Reddit Comments. In Proceedings of the 59th Annual Meeting of the Association of Computational Linguistics (ACL-2021), August 2021.

12. Saif M. Mohammad. 2020. Examining Citations of Natural Language Processing Literature. In Proceedings of the 58th Annual Meeting of the Association of Computational Linguistics (ACL-2020), Seattle, USA.

13. Saif M. Mohammad. 2020. Gender Gap in Natural Language Processing Research: Disparities in Authorship and Citations. In Proceedings of the 58th Annual Meeting of the Association of Computational Linguistics (ACL-2020), Seattle, USA.

14. Chris K. Kim, Uta Hinrichs, and Saif M. Mohammad, and Chris Collins. 2020. Lexichrome: Text Construction and Lexical Discovery with Word-Color Associations Using Interactive Visualization. In Proceedings of DIS 2020: Conference on Designing Interactive Systems, Eindhoven, Netherlands.

15. Saif M. Mohammad. 2020. NLP Scholar: A Dataset for Examining the State of NLP Research. In Proceedings of the 12th Language Resources and Evaluation Conference (LREC-2020), Marseille, France.

16. Saif M. Mohammad. 2020. NLP Scholar: An Interactive Visual Explorer for Natural Language Processing Literature. In Proceedings of the 58th Annual Meeting of the Association of Computational Linguistics (ACL-2020), Seattle, USA.

17. Will E. Hipson, and Saif M. Mohammad. 2020. PoKi: A Large Dataset of Poems by Children. In Proceedings of the 12th Language Resources and Evaluation Conference (LREC-2020), Marseille, France.

18. Svetlana Kiritchenko, Will Hipson, Robert Coplan, and Saif M. Mohammad. 2020. SOLO: A Corpus of Tweets for Examining the State of Being Alone. In Proceedings of the 12th Language Resources and Evaluation Conference (LREC-2020), Marseille, France.

19. Saif M. Mohammad. 2020. WordWars: A Dataset to Examine the Natural Selection of Words. In Proceedings of the 12th Language Resources and Evaluation Conference (LREC-2020), Marseille, France.

20. Shima Asaadi, Saif M. Mohammad, and Svetlana Kiritchenko. 2019. Big Bird: A Large, Fine-Grained, Bigram Relatedness Dataset for Examining Semantic Composition. In *Proceedings of*

the North American Chapter of the Association for Computational Linguistics (NAACL-2019), Minneapolis, MN, USA.

21. Michael Wojatzki, Torsten Zesch, Saif M. Mohammad, and Svetlana Kiritchenko. 2018. Agree or Disagree: Predicting Judgments on Nuanced Assertions. In Proceedings of *Sem, New Orleans, LA, USA.

22. Svetlana Kiritchenko and Saif M. Mohammad. 2018. Examining Gender and Race Bias in Two Hundred Sentiment Analysis Systems. In Proceedings of *Sem, New Orleans, LA, USA.

23. Saif M. Mohammad. 2018. Obtaining Reliable Human Ratings of Valence, Arousal, and Dominance for 20,000 English Words. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, Melbourne, Australia.

24. Michael Wojatzki, Saif M. Mohammad, Torsten Zesch, and Svetlana Kiritchenko. 2018. Quantifying Qualitative Data for Understanding Controversial Issues. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.

25. Saif M. Mohammad and Svetlana Kiritchenko. 2018. Understanding Emotions: A Dataset of Tweets to Study Interactions between Affect Categories. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.

26. Saif M. Mohammad and Svetlana Kiritchenko. 2018. WikiArt Emotions: An Annotated Dataset of Emotions Evoked by Art. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.

27. Saif M. Mohammad. 2018. Word Affect Intensities. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.

28. Svetlana Kiritchenko and Saif M. Mohammad. 2017. Best-Worst Scaling More Reliable than Rating Scales: A Case Study on Sentiment Intensity Annotation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL-2017)*, Vancouver, Canada.

29. Saif M. Mohammad and Felipe Bravo-Marquez. 2017. Emotion Intensities in Tweets. In *Proceedings of the Sixth Joint Conference on Lexical and Computational Semantics (*Sem)*, Vancouver, Canada.

30. Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. A Dataset for Detecting Stance in Tweets. In *Proceedings of the 10th edition of the Language Resources and Evaluation Conference*, Portorož, Slovenia.

31. Svetlana Kiritchenko and Saif M. Mohammad. 2016. Capturing Reliable Fine-Grained Sentiment Associations by Crowdsourcing and Best–Worst Scaling. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, San Diego, CA.

32. Parinaz Sobhani, Saif M. Mohammad, and Svetlana Kiritchenko. 2016. Detecting Stance in Tweets And Analyzing its Interaction with Sentiment. In *Proceedings of the Joint Conference on Lexical and Computational Semantics (*Sem)*, Berlin, Germany.

33. Felipe Bravo-Marquez, Eibe Frank, Saif M. Mohammad, and Bernhard Pfahringer. 2016. Determining Word-Emotion Associations from Tweets by Multi-Label Classification. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence (WI'16)*, Omaha, Nebraska, USA.

34. Svetlana Kiritchenko and Saif M. Mohammad. 2016. Happy Accident: A Sentiment Composition Lexicon for Opposing Polarity Phrases. In *Proceedings of the 10th edition of the Language Resources and Evaluation Conference*, Portorož, Slovenia.

35. Saif M. Mohammad, Ekaterina Shutova, and Peter Turney. 2016. Metaphor as a Medium for Emotion: An Empirical Study. In *Proceedings of the Joint Conference on Lexical and Computational Semantics (*Sem)*, Berlin, Germany.

36. Svetlana Kiritchenko and Saif M. Mohammad. 2016. Sentiment Composition of Words with Opposing Polarities. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, San Diego, CA.

37. Saif M. Mohammad, Mohammad Salameh, and Svetlana Kiritchenko. 2016. Sentiment Lexicons for Arabic Social Media. In *Proceedings of the 10th edition of the Language Resources and Evaluation Conference*, Portorož, Slovenia.

38. Mohammad Salameh, Saif M. Mohammad and Svetlana Kiritchenko. 2015. Sentiment After Translation: A Case-Study on Arabic Social Media Posts. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL-2015)*, Denver, Colorado.

39. Xiaodan Zhu, Hongyu Guo, Saif M. Mohammad and Svetlana Kiritchenko. 2014. An Empirical Study on the Effect of Negation Words on Sentiment. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, Baltimore, MD.

40. Saif M. Mohammad. 2012. #Emotional Tweets. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics (*Sem)*, Montreal, Canada.

41. Saif M. Mohammad. 2012. Portable Features for Classifying Emotional Text. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Montreal, Canada.

42. Saif M. Mohammad. 2011. Even the Abstract have Colour: Consensus in Word-Colour Associations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-2011)*, Portland, OR.

43. Yuval Marton, Saif M. Mohammad, and Philip Resnik. 2009. Estimating semantic distance using soft semantic constraints in knowledge-source–corpus hybrid models. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2009)*, Singapore.

44. Saif M. Mohammad, Bonnie Dorr, and Cody Dunne. 2009. Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2009)*, Singapore.

45. Saif M. Mohammad, Bonnie Dorr, Melissa Egan, Ahmed Hassan, Pradeep Muthukrishan, Vahed Qazvinian, Dragomir Radev, and David Zajic. 2009. Using citations to generate surveys of scientific paradigms. In *Proceedings of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT-2009)*, Boulder, Colorado.

46. Saif M. Mohammad, Bonnie Dorr, and Graeme Hirst. 2008. Computing word-pair antonymy. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2008)*, Waikiki, Hawaii.

47. Saif M. Mohammad, Iryna Gurevych, Graeme Hirst, and Torsten Zesch. 2007. Cross-lingual distributional profiles of concepts for measuring semantic distance. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007),*

48. Saif M. Mohammad and Graeme Hirst. 2006. Determining word sense dominance using a thesaurus. In *Proceedings of the Conference of the European chapter of the Association for Computational Linguistics (EACL-2006)*, Trento, Italy.

49. Saif M. Mohammad and Graeme Hirst. 2006. Distributional measures of concept-distance: A task-oriented evaluation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2006),* Sydney, Australia.

50. Saif M. Mohammad and Ted Pedersen. 2004. Combining lexical and syntactic features for supervised word sense disambiguation. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL-2004)*, Boston, MA.

51. Saif M. Mohammad and Ted Pedersen. 2003. Guaranteed pre-tagging for the Brill Tagger. In *Proceedings of the International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2003)*, Mexico City, Mexico.

Refereed Workshop Papers

1. Shamsuddeen Hassan Muhammad, Idris Abdulmumin, Abinew Ali Ayele, Nedjma Ousidhoum, David Ifeoluwa Adelani, Seid Muhie Yimam, Ibrahim Sa'id Ahmad, Meriem Beloucif, Saif Mohammad, Sebastian Ruder, Oumaima Hourrane, Pavel Brazdil, Felermino Dário Mário António Ali, Davis Davis, Salomey Osei, Bello Shehu Bello, Falalu Ibrahim, Tajuddeen Gwadabe, Samuel Rutunda, Tadesse Belay, Wendimu Baye Messelle, Hailu Beshada Balcha, Sisay Adugna Chala, Hagos Tesfahun Gebremichael, Bernard Opoku, Steven Arthur. AfriSenti: A Twitter Sentiment Analysis Benchmark for African Languages. In *Proceedings of the AfricaNLP workshop at ICLR-2023*, Kigali, Rwanda. [Best Paper Award]

2. Shamsuddeen Hassan Muhammad, Idris Abdulmumin, Seid Muhie Yimam, David Ifeoluwa Adelani, Ibrahim Sa'id Ahmad, Nedjma Ousidhoum, Abinew Ayele, Saif M. Mohammad, Meriem Beloucif, Sebastian Ruder. SemEval-2023 Task 12: Sentiment Analysis for African Languages (AfriSenti-SemEval). In Proceedings of the ACL 2023 workshop SemEval 2023, Toronto, Canada.

3. Daniela Teodorescu, Alona Fyshe, Saif M. Mohammad. Utterance Emotion Dynamics in Children's Poems: Emotional Changes Across Age. In *Proceedings of the ACL 2023 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, Toronto, Canada.

4. Katherine Fraser, Frauke Zeller, David Harris Smith, Saif M. Mohammad, and Frank Rudicz. 2019. How do we feel when a robot dies? Emotions expressed on Twitter before and after hitchBOT's destruction. In *Proceedings of the NAACL 2019 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, Minneapolis, MN.

5. Habibeh Naderi, Svetlana Kiritchenko, Saif M. Mohammad, and Stan Matwin. 2018. DeepMiner at SemEval-2018 Task 1: Emotion Intensity Recognition Using Deep Representation Learning. In Proceedings of International Workshop on Semantic Evaluation (SemEval-2018), New Orleans, LA, USA.

6. Roman Klinger, Orphée De Clercq, Saif M. Mohammad, and Alexandra Balahur. 2018. IEST: WASSA-2018 Implicit Emotions Shared Task. In *Proceedings of the EMNLP Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, Brussels, Belgium.

7. Saif M. Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. Semeval-2018 Task 1: Affect in tweets. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA.

8. Svetlana Kiritchenko, Saif M. Mohammad, Jason Morin, and Berry de Bruijn. 2017. NRC-Canada at SMM4H Shared Task: Classifying Tweets Mentioning Adverse Drug Reactions and Medication Intake. In *Proceedings of the Social Media Mining for Health Applications Workshop at AMIA-2017*, Washington, DC, USA.

9. Saif M. Mohammad and Felipe Bravo-Marquez. 2017. WASSA-2017 Shared Task on Emotion Intensity. In *Proceedings of the EMNLP 2017 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, Copenhagen, Denmark.

10. Saif M. Mohammad. 2016. A Practical Guide to Sentiment Annotation: Challenges and Solutions. In *Proceedings of the NAACL 2016 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, San Diego, California.

11. Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. Semeval-2016 Task 6: Detecting Stance in Tweets. In *Proceedings of the International Workshop on Semantic Evaluation (SemEval-2016)*, San Diego, CA.

12. Svetlana Kiritchenko, Saif M. Mohammad, and Mohammad Salameh. 2016. Semeval-2016 Task 7: Determining Sentiment Intensity of English and Arabic Phrases. In *Proceedings of the International Workshop on Semantic Evaluation (SemEval-2016)*, San Diego, CA.

13. Svetlana Kiritchenko and Saif M. Mohammad. 2016. The Effect of Negators, Modals, and Degree Adverbs on Sentiment Composition. In *Proceedings of the NAACL 2016 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media*, San Diego, CA.

14. Saif M. Mohammad. 2015. Imagisaurus: An Interactive Visualizer of Valence and Emotion in the Roget's Thesaurus. In *Proceedings of the EMNLP 2015 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, Lisbon, Portugal.

15. Sara Rosenthal, Preslav Nakov, Svetlana Kiritchenko, Saif M Mohammad, Alan Ritter, and Veselin Stoyanov. 2015. SemEval-2015 Task 10: Sentiment Analysis in Twitter. In *Proceedings of the international workshop on Semantic Evaluation Exercises (SemEval 2015)*, Denver, CO.

16. Hannah Davis and Saif M. Mohammad. 2014. Generating Music from Literature. In *Proceedings of the EACL Workshop on Computational Linguistics for Literature*, Gothenburg, Sweden.

17. Svetlana Kiritchenko, Xiaodan Zhu, Colin Cherry, and Saif M. Mohammad. 2014. NRC-Canada-2014: Detecting Aspects and Sentiment in Customer Reviews. In *Proceedings of the international workshop on Semantic Evaluation Exercises (SemEval-2014)*, Dublin, Ireland.

18. Xiaodan Zhu, Svetlana Kiritchenko, and Saif M. Mohammad. 2014. NRC-Canada-2014: Recent Improvements in the Sentiment Analysis of Tweets. In *Proceedings of the eighth international workshop on Semantic Evaluation Exercises (SemEval-2014)*, Dublin, Ireland.

19. Saif M. Mohammad, Xiaodan Zhu, and Joel Martin. 2014. Semantic Role Labeling of Emotions in Tweets. In *Proceedings of the ACL 2014 Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media (WASSA)*, Baltimore, MD.

20. Saif M. Mohammad, Svetlana Kiritchenko and Joel Martin. 2013. Identifying Purpose Behind Electoral Tweets. In *Proceedings of the KDD Workshop on Issues of Sentiment Discovery and Opinion Mining (WISDOM-2013)*, Chicago, USA.

21. 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, USA.

22. Saif M. Mohammad and Svetlana Kiritchenko. 2013. Using Nuances of Emotion to Identify Personality. In *Proceedings of the ICWSM Workshop on Computational Personality Recognition*, Boston, USA.

23. David Jurgens, Saif M. Mohammad, Peter Turney, and Keith Holyoak. 2012. SemEval-2012 Task 2: Measuring Degrees of Relational Similarity. In *Proceedings of the sixth international workshop on Semantic Evaluation Exercises (SemEval-2012)*, Montreal, Canada.

24. Saif M. Mohammad. 2011. Colourful Language: Measuring Word-Colour Associations. *Proceedings of the ACL 2011 Workshop on Cognitive Modeling and Computational Linguistics (CMCL)*, Portland, OR.

25. Saif M. Mohammad. 2011. From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales. *Proceedings of the ACL 2011 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH)*, Portland, OR.

26. Alistair Kennedy, Anna Kazantseva, Saif M. Mohammad, Terry Copeck, Diana Inkpen, and Stan Szpakowicz. 2011. Getting Emotional About News. Proceedings of the Text Analysis Conference (TAC-2011), Gaithersburg, MD.

27. Saif M. Mohammad and Tony Yang. 2011. Tracking Sentiment in Mail: How Genders Differ on Emotional Axes. *Proceedings of the ACL 2011 Workshop on ACL 2011 Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA)*, Portland, OR.

28. Saif M. Mohammad and Peter Turney. 2010. Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. *Proceedings of the NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, LA, California.

29. James Mayfield, Bonnie Dorr, Jason Eisner, Tim Finin, Saif M. Mohammad, Douglas Oard, Ralph Weischedel, David Yarowsky, and others. 2009. Cross-Document Coreference Resolution: A Key Technology for Learning by Reading. In *Proceedings of the AAAI Spring Symposium on Learning by Reading and Learning to Read (AAAI-09)*, Menlo Park, CA.

30. Saif M. Mohammad, Bonnie Dorr, Melissa Egan, Jimmy Lin, and David Zajic. 2008. Multiple alternative sentence compressions and word-pair antonymy for automatic text summarization and recognizing textual entailment. *Proceedings of the Text Analysis Conference (TAC-2008)*, Gaithersburg, MD.

31. Saif M. Mohammad, Bonnie Dorr, and Graeme Hirst. 2008. Towards antonymy-aware natural language applications. *Proceedings of the Symposium on Semantic Knowledge Discovery, Organization and Use (SKDOU-2008)*, New York, NY.

32. Saif M. Mohammad, Graeme Hirst, and Philip Resnik. 2007. TOR, TORMD: Distributional profiles of concepts for unsupervised word sense disambiguation. *Proceedings of the International Workshop on Semantic Evaluations (SemEval-2007)*, Prague, Czech Republic.

33. Saif M. Mohammad and Ted Pedersen. 2004. Complementarity of lexical and simple syntactic features: The SyntaLex approach to SensEval-3. *Proceedings of the International Workshop on Semantic Evaluations (SensEval-3)*, Barcelona, Spain.

Theses

1. Saif M. Mohammad. 2008. Measuring semantic distance using distributional profiles of concepts. Ph.D. thesis, Department of Computer Science. University of Toronto, Canada.

2. Saif M. Mohammad and Ted Pedersen. 2004. Combining lexical and syntactic features for supervised word sense disambiguation. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL-2004)*, Boston, MA.

Technical Reports

1. Saif M. Mohammad. Ethics Sheets for AI Tasks. arXiv preprint arXiv:2107.01183. July 2021.

2. Will E. Hipson and Saif M. Mohammad. Emotion Dynamics in Movie Dialogues. arXiv preprint arXiv:2103.01345. March 2021.

3. Saif M. Mohammad. Practical and Ethical Considerations in the Effective use of Emotion and Sentiment Lexicons. arXiv preprint arXiv:2011.03492. December 2020.

4. Berry de Bruijn, Alain Désillets, Katherine Fraser, Svetlana Kiritchenko, Saif M. Mohammad, Norm Vinson, Peter Bloomfield, Hayley Brace, Katarzyna Brzoska, Anat Elhalal, K. Ho, Libby Kinsey, Ross McWhirter, Madalina Nazare, and Evelyn Ofuri-Kuragu. 2019. Applied AI Ethics: Report on Canada-United Kingdom Symposia on Ethics in AI in Ottawa, Canada and London, UK. Digital Catapult, London, UK / NRC, Ottawa, Canada.

5. Saif M. Mohammad. 2019. The State of NLP Literature: A Diachronic Analysis of the ACL Anthology. arXiv preprint arXiv:1911.03562.

6. Saif M. Mohammad and Peter Turney. 2013. NRC Emotion Lexicon. Technical Report, National Research Council Canada, Ottawa, Canada.

7. Saif M. Mohammad, Bonnie J. Dorr, Graeme Hirst, and Peter D. Turney. 2011. Measuring Degrees of Semantic Opposition. Technical Report, NRC, Ottawa, Canada.

8. Saif M. Mohammad. 2011. Sentiment Analysis of Mail and Books. Technical Report, National Research Council Canada, Ottawa, Canada.

9. Saif M. Mohammad and Graeme Hirst. 2007. Distributional Measures of Semantic Distance: A Survey. arXiv:1203.1858.

10. Saif M. Mohammad and Graeme Hirst. 2006. Distributional Measures as Proxies for Semantic Relatedness. arXiv:1203.1889.

Symposia

1. Gregory J. Park, Saif M. Mohammad, and Johannes C. Eichstaedt. 2015. My N is Ten Million: Using Social Media to Track Emotion, Mental Health, and Measure Personality Across Entire Populations. In *Proceedings of the Symposium at the International Convention of Psychological Science (ICPS)*, Amsterdam, The Netherlands.

2. Bonnie Dorr, Saif M. Mohammad, and Boyan Onyshkevych. 2008. From linguistic annotations to knowledge objects. *Proceedings of the Symposium on Semantic Knowledge Discovery, Organization and Use (SKDOU-2008)*, New York, NY.

Scientific Blog Posts

1. Ethics Sheets for AI Tasks. July 5, 2021.
2. Ethics Sheet for Automatic Emotion Recognition and Sentiment Analysis. July 5, 2021.
3. Ten Years of the NRC Word-Emotion Association Lexicon. March, 2020.
4. The State of NLP Literature: A Diachronic Analysis of the ACL Anthology. October, 2019.

DATASETS CREATED

Manually Labeled Datasets

Year	Project	# Datasets
2023	A Twitter Sentiment Analysis Benchmark for African Languages	14
2023	BIG-Bench: Beyond the Imitation Game Benchmark Dataset for over 200 Tasks	1
2021	Reddit comments annotated for degree of supportiveness and offensiveness	1
2018	Art annotated for emotion, likability, and more	3
2018	SemEval-2018 Task 1: Affect in Tweets Data (5 tasks, 3 languages, 5 affect dim.)	31
2018	Equity Evaluation Corpus (for Gender and Race Bias)	2
2018	Assertions in Controversial Issues	2
2018	Valence, Arousal, and Dominance Lexicon	3
2018	Word Affect Intensities Lexicon	8
2017	Emotion Intensities in Tweets Dataset	4
2016	Stance through Tweets Dataset	2
2016	Metaphor as a Medium for Emotion Dataset	2
2016	2016 Task 7: Sentiment Intensity of English and Arabic Phrases	2
2016	Arabic Sentiment Lexicons	3
2016	Sentiment Composition Lexicon of Negators, Modals, and Adverbs	1
2016	Sentiment Composition Lexicon of Opposing Polarity Phrases	1
2015	SemEval-2015 Task 10: Sentiment Analysis in Twitter	5
2015	Sentiment in Arabic BBN Blog Posts and Syrian Tweets	2
2015	SemEval-2015 English Twitter Sentiment Lexicon	1
2014	Electoral Tweets Dataset	4
2012	Word–Word Relational Similarity Dataset	1
2010	NRC Emotion Lexicon	10
2010	Word–Color Association Lexicon	1
Total		**104**

Automatically Compiled Datasets

Year	Project	# Datasets
2023	A Dataset of ∼2.5B Citations from and to NLP Papers	1
2023	A dataset of 70K NLP Papers and their Citation Networks Over Time	1
2022	D3: A Dataset of Metadata for ∼81M Computer Science Research Papers	1
2020	Tweets from Canadian and US Cities during the 2020 Pandemic	2
2020	PoKi: A Dataset of 72K Poems Written by Children	1
2020	SOLO: A Dataset of tweets that mention solitude and loneliness	1
2020	NLP Scholar: A dataset of 50K NLP papers and their citations	1
2020	WordWars: A Dataset to Examine the Natural Selection of Words	1
2019	HitchBot: A Dataset of Tweets mentioning the Canadian hitchhiking robot	1
2015	Tweets with emotion word hashtags	1
2013	NRC Hashtag Emotion Lexicon	1
2013	NRC Hashtag Sentiment Lexicons	3
2013	Sentiment140 Regular, Affirmative Context, and Negated Context Lexicons	3
2013	Yelp and Amazon Sentiment Lexica	2
2013	Lexical Contrast Datasets	4
2012	#Emotional Tweets	8
Total		**32**

NOTABLE PRESS MENTIONS

1. CBC Radio, Feb 12, 2021. Automatic emotion detection: benefits and potential harms.
2. The Economist, Jan 12, 2021. How Donald Trump evolved into a prolific, angry Twitter user. Use of the NRC Emotion lexicon to Analyze emotions in tweets by Donald Trump.
3. AdNews, April 20, 2020. Coronavirus: Analysing stories from the Twitter frontline. Use of the NRC Emotion lexicon to Analyze emotions in tweets during the Covid-19 pandemic.
4. Tech Xplore, February 26, 2020. A language generation system that can compose creative poetry. The NRC Emotion Lexicon is used along with a deep learning system for automatic poetry generation. Original paper: Introducing Aspects of Creativity in Automatic Poetry Generation.
5. Universo Online, Dec 5, 2019. Gender bias in sentiment analysis systems and gender disparities in language.
6. Politico.mx, August 28, 2019. Enojo o sorpresa: emociones en Informes de López Portillo a Calderó. The NRC Emotion Lexicon is used to analyze the speeches of the Mexican President and other politicians.
7. Gender and Racial Bias in Cloud NLP Sentiment APIs, Aug 21, 2019. Article looking into race and gender biases in the Google and AWS cloud sentiment analysis APIs using the Equity Evaluation Corpus and the techniques we published in 2018.
8. The New York Times, February 26, 2019. Extreme Weather Can Feel 'Normal' After Just a Few Years, Study Finds. Brief interview on my opinion on work on perceptions of climate change.
9. Psychology Today, December 27, 2018. Emotional Contagions Can Spread Like Wildfire Via YouTube. The NRC Emotion Lexicon is used to analyze the emotions in user comments on vlogs. Original Paper: Multilevel emotion transfer on YouTube: Disentangling the effects of emotional contagion and homophily on video audiences.
10. InformationWeek, November 15, 2018. Decoding Programmers: How Emotions Can Change Code. Analyzing developer comments to show that their emotions impact their code.
11. Medium, DAIR, Artificial Intelligence Research, Perspectives, and Technologies, Aug 2, 2018. Examining Gender and Race Bias in Sentiment Analysis Systems.
12. SeekingAlpha, Apr. 30, 2018. A Sentiment Analysis Approach To Predicting Stock Returns. The NRC Emotion Lexicon is used to analyze 10-K reports of S&P 500 companies to predict future stock returns.
13. Technical.ly, February 1, 2018. Who's the angriest character on 'Seinfeld'? The NRC Emotion Lexicon is used to analyze the language of the characters in Seinfeld.

14. January 24, 2018. An examination of the Twitter habits of Kentucky lawmakers. The NRC Emotion Lexicon is used to analyze the tweets of Democrats and Republicans.

15. April 13, 2017. Faster but less furious: why the Fast and Furious franchise is a box office goldmine. The NRC Emotion Lexicon is used to analyze the emotions in the Fast and Furious series movie scripts.

16. The Crosstab, March 2, 2017. Trump's SOTU vs. the Past — Sentiment Analysis and Topic Modeling.

17. Esquire, February 23, 2017. All Radiohead Songs Are Sad, but This Graph Shows Which Are the Saddest. The NRC Emotion Lexicon is used to compute the degree of sadness for Radiohead songs. Interactive Chart. Also picked up in: A Journal of Musical Things, February 24, 2017. Science Discovers Which Radiohead Song is the Saddest

18. Thomas Park, February 28. Where I End and You Begin: Finding the Most Depressing Radiohead Songs using Crowd Data from SongMeanings.

19. Religion in Public, February 13, 2017. Evangelical leaders are as negative on immigration and Islam as Donald Trump.

20. February 11, 2017. Text Analysis of NHL Hockey Coach Interviews.

21. Washington Post, October 22, 2016. Donald Trump and Hillary Clinton took to the debate stage and made sweet, sweet music. The article mentions NRC in relation to the lexicons we created, which were used to generate music from the Trump-Hillary debate text.

22. Washington Post, CBS News, Columbia Tribune, and others, September 23, 2016. This symphony had both human and computer composers. Articles about a symphony orchestra performed music composed using the NRC Emotion Lexicon under the glass of the Louvre museum in Paris on Sept. 20, 2016.

23. Washington Post, August 12, 2016. Two people write Trump's tweets. He writes the angrier ones.

24. BGR, August 11, 2016. Donald Trump's angriest tweets are sent from his Android while the nice ones are sent from an iPhone.

25. NYC Data Science Academy, August 7, 2016. Twitter Analysis of Presidential Candidates 2016.

26. Variance Explained, August 6, 2016. Text analysis of Trump's tweets confirms he writes only the (angrier) Android half. News also picked up by NPR, Los Angeles Times, Scientific American, The Verge, and others.

27. The Telegraph, June 15, 2016: EU referendum: Remain uses Project Fear more in tweets than Leave, analysis shows. [Use of the NRC Emotion Lexicon, aka EmoLex, to track sentiment in EU referendum tweets (Brexit).]

28. An art project, the Wishing Wall, that uses the NRC Emotion lexicon was displayed in:

    - Tekniska Museet, Stockholm, Sweden (Oct 14 – Aug 15)

- Onassis Cultural Centre, Athens (19th Oct'15 – 10th Jan'16)
- Zorlu Centre in Istanbul (16th Feb – 12th June'16)

29. Fast Company, March 25, 2016: An Emotional Map Of The City, As Captured Through Its Sounds. [Use of the NRC Emotion Lexicon, aka EmoLex to create Chatty Maps.]
30. Press release stating the use of EmoLex to create Chatty Maps.
31. PC World, May 15, 2014: AI System Reads Novels, Writes Music for Them.
32. Popular Science, May 14, 2014: Robot Reads Novels, Writes Songs about Them.
33. io9, May 12, 2014: Researchers Train Computers to Manipulate Human Emotions with Art.
34. LiveScience, May 11, 2014: 'TransProse' Software Creates Musical Soundtracks from Books.
35. TIME, May 7, 2014: This Is What Classic Novels Sound Like When a Computer Turns Them Into Piano Music.
36. SlashDot, March 23, 2014: Algorithm Composes Music By Text Analyzing the World's Best Novels.
37. The Physics arXiv Blog, March 20, 2014: The Music Composed By An Algorithm Analysing The World's Best Novels.
38. Glass Hammer, December 3, 2013: Are Your Emails Communicating a Lack of Confidence?
39. Singularity Hub, November 10, 2013: Algorithm Tracks Literary Emotion in Shakespeare, the Brothers Grimm.
40. The Physics ArXiv, October 4, 2013: Data Mining Reveals the Emotional Differences in Emails Written by Men and Women.
41. SlashDot, October 4, 2013: Data Mining Reveals the Emotional Differences In Emails From Men and Women.
42. The Physics ArXiv, October 1, 2013: Text Analyser Reveals Emotional Temperature of Novels and Fairy Tales. Also in EduBits.
43. SlashDot, October 1, 2013: Text Analyzer Reveals Emotional 'Temperature' of Novels and Fairy Tales.
44. The New Scientist, September 27: What your email style says about your personality. Also in Times of India, MSN, Pharmacon, Galileo, Amic, and others.
45. Article in MIT Technology Review, September 5, 2013: How Mechanical Turkers Crowdsourced a Huge Lexicon of Links Between Words and Emotion.
46. TIME, August 14, 2013: Main Tweet: Researchers Dig Into The Intersection of Politics and Twitter.