

The Dark Side of NLP: Gefahren automatischer Sprachverarbeitung

Michael Strube

michael.strube at h-its.org

July 30, 2018

Inhalt

Computerlinguistische Anwendungen haben sich im Alltag durchgesetzt: Suchmaschinen, Rechtschreibkorrektur, maschinelle Übersetzung, Spracherkennung usw. stehen jedem auf Computer und Mobiltelefon zur Verfügung. Die Computerlinguistik hilft allerdings nicht nur uns Endanwendern, sondern auch dem Handel, mehr über seine Kunden zu erfahren, der Industrie, personalisierte Werbung zu plazieren, autoritären Staaten, Microblogeinträge zu zensieren, Geheimdiensten, Telefongespräche und Emails nicht mehr nur auf Stichwörter sondern auch auf Inhalte hin zu durchsuchen. Im Seminar sollen Methoden und Anwendungen aus den Bereichen Sentiment Analysis, maschinelle Übersetzung, Textmining, NLP und Social Media, Computational Advertising usw. untersucht und im Hinblick auf ihre gesellschaftlichen und ethischen Auswirkungen hinterfragt werden, etwa:

- Wo ist die Grenze zwischen "cooler" und gefährlicher Forschung?
- Sind wir gerade dabei die Atomphysiker des 21. Jahrhunderts zu werden?
- Wie können wir uns auf den durch Fortschritte in den Bereich Machine und Deep Learning vorhergesagten gesellschaftlichen Wandel vorbereiten?
- Wieso beschäftigen sich unsere besten Köpfe mit Werbung?
- Und: Was können wir tun?

Das Thema verlangt, das Seminar anders zu durchzuführen als üblich (Seminarvortrag plus mühsame Diskussion). Wir werden andere Formen der Interaktion erproben und einüben.

Literatur

- Blue, Ethan et al. (2014). Engineering and War: Militarism, Ethics, Institutions, Alternatives. Morgan and Claypool Publishers.

Artikel

- Hovy, Dirk and Spruit, Shannon L. (2016). The social impact of Natural Language Processing. In Proc. of ACL-16 Short Papers, pp.591-598.

Termine, Themenvorschläge

08.11.2016

Einführung, Motivation, ...

15.11.2016

Michael Strube

Geschichte – 2. Weltkrieg, Vietnam, Förderung für KI und CL durch das Militär, Snowden, Metadaten, Datenschutz:

(Shapley, 1972; Glantz & Albers, 1974; Holden, 1975; Edsall, 1975; Thompson, 1986; Schuler & Jacky, 1989; Beusmans & Wieckert, 1989; Winograd, 1991; Yen, 2004; Popp et al., 2004; Coffman et al., 2004; Bradford, 2006; Rubenstein et al., 2008; Cho, 2013; Hajaj et al., 2013; Greenwald, 2014; Schneier, 2014; 2015)

zur Vorbereitung: Stellungnahme zum Fall Snowden

22.11.2016

Michael Strube

Wie funktioniert Google?

(?; Schmidt & Cohen, 2013)

Attila Azgin

Wie funktioniert Facebook?

zur Vorbereitung: (nicht notwendigerweise persönliche) Stellungnahme zu (a) “Weshalb können wir ohne Google nicht leben? Ist es die Qualität der Suchmaschine, die Bequemlichkeit all der anderen Dienste oder ...?” oder (b) “Weshalb sind wir bei Facebook? Sozialer Druck, Exhibitionismus oder ...?”

29.11.2016

David Stronczek

Microblogs – Benutzer, Entities, Events:

(Benton et al., 2016; Kairam et al., 2013)

Denis Orechov

Microblogs – *Hate Speech* Erkennung, Dogmatismus:

(Waseem & Hovy, 2016; Waseem, 2016; Mehdad & Tetreault, 2016; Fast & Horvitz, 2016)

zur Vorbereitung (jeweils zwei Fragen zu):

(Benton et al., 2016) und (Warner & Hirschberg, 2012)

optional

Microblogs – Event Extraction:

(Benson et al., 2011; Zhou et al., 2016a)

optional

Microblogs – Entity Linking:

(Guo et al., 2013b; Liu et al., 2013; Dredze et al., 2016; Yang et al., 2016; Peng & Dredze, 2016; Strauss et al., 2016) und Entity Linking allgemein: (Ji & Grishman, 2011; Ji et al., 2016)

auch: <http://www.akbc.ws/2014/slides/onyshkevych-nips-akbc.pdf>

optional

Microblogs – Inhaltserschließung:

(Diao & Jiang, 2013; Guo et al., 2013a; Eisenstein, 2013; Chua & Asur, 2013; Grinberg et al., 2013; Tsur & Rappoport, 2013; Zhang et al., 2016b; Blodgett et al., 2016)

06.12.2016 **geänderte Zeit: 16.00-17.30**

wegen Vortrag zu “Gegenwartsliteratur und Zeitdiagnostik: Dave Eggers’ Roman ’The Circle’. Segen und Fluch der Transparenzgesellschaft” – 18.15, Hörsaal 14, Neue Universität

Catarina Cramer, Maximilian Müller-Eberstein

Microblogs – Soziale Faktoren, Profiling:

(Poulston et al., 2016; Volkova & Bachrach, 2016) und evtl. (Flekova et al., 2016b)

zur Vorbereitung: (Fragen einsenden bis spätestens Montag 13 Uhr)

(Poulston et al., 2016) und (Volkova & Bachrach, 2016)

optional

(Argamon et al., 2009; Burger et al., 2011; Bamman et al., 2012a; Ciot et al., 2013; Eisenstein et al., 2011; Rangel et al., 2013; Nguyen et al., 2013a; Bergsma & Van Durme, 2013; Flekova & Gurevych, 2013; Flekova et al., 2016a; 2016b; Kim et al., 2016; ?; Liu et al., 2016a; Kamijo et al., 2016; Xiao et al., 2016)

13.12.2016

fällt aus: Michael Strube at COLING 2016

20.12.2016

Dominik Both

Microblogs – Autorenerkennung:

(Sousa Silva et al., 2011; Wang et al., 2013; Schwartz et al., 2013; Qian & Liu, 2013; Rocha et al., 2017)

Anti-Profiling:

Raphael Schumann

(Brennan et al., 2012; Reddy & Knight, 2016)

10.01.2017

Theresa Sick und Aaaaaa Aaaaaaa Aaaaaa

Soziale Netzwerke und NLP – soziologische, psychologische Pänomene:

(Sleeper et al., 2013; Nitta et al., 2013; Ptaszynski et al., 2016)

zur Vorbereitung:

(Sleeper et al., 2013) und ((Nitta et al., 2013) oder (Ptaszynski et al., 2016))

optional

(Elson et al., 2010; Rao et al., 2011; Cano et al., 2013; Abu-Jbara et al., 2013; Agarwal et al., 2013; ?; El-Arini et al., 2013; Wohlgenannt et al., 2016; Litvak et al., 2016)

17.01.2017

Julian Rodriquez

Soziale Netzwerke – Facebook:

“The Facebook Study” (Kramer, 2012; Kramer et al., 2014)

optional: (Das & Kramer, 2013; Verma, 2014), s. auch: <https://research.facebook.com/>

Julian Hitschler

Ethik und Machine Learning:

(Bostrom, 2014, nur die ersten zwei Abschnitte), (Burri, 2016) und <https://medium.com/@hannawallach/big-data-machine-learning-and-the-social-sciences-927a8e20460d#.s2956x22r>

<https://medium.com/@hannawallach/big-data-machine-learning-and-the-social-sciences-927a8e20460d#.s2956x22r>

Ines Pisetta

Psychologie und NLP – Erkennung von Lügen:

(Ott et al., 2011)

optional: (Burgoon et al., 2003; Zhou et al., 2003; 2004; Bond & Lee, 2005; Hancock et al., 2005; Graciarena et al., 2006; Bachenko et al., 2008; Takase et al., 2013; Feng & Hirst, 2013; Li et al., 2013a; Resnik et al., 2013; Li et al., 2013b; 2013c; Ott et al., 2013)

zur Vorbereitung (2 aus den folgenden 3 (4)): (Ott et al., 2011) und/oder (Kramer et al., 2014), und/oder ((Burri, 2016) oder <https://medium.com/@hannawallach/big-data-machine-learning-and-the-social-sciences-927a8e20460d#.s2956x22r>)

Zusatzfragen für die letzte Sitzung!

<http://www.cl.uni-heidelberg.de/courses/ws16/darkside/material/zusatzfragen.txt>

24.01.2017

Psychologie und NLP – Allgemein:

(Coppersmith et al., 2016)

Julia Suter, Till Nocher

Psychologie und NLP – Erkennung von Depression etc.:

(Le et al., 2011; Mowery et al., 2016)

optional: (Stirman & Pennebaker, 2001; Lott et al., 2002; Rude et al., 2004; Cohn et al., 2004; Pestian et al., 2012; Resnik et al., 2013; Lamb et al., 2013; De Choudhury et al., 2013; Nguyen et al., 2013b; Coppersmith et al., 2016; Sadeque et al., 2016; Smith et al., 2016; Liu et al., 2016b)

Oliver Singler, Sebastian Burst

Psychologie und NLP – Macht, Einfluß, etc.

(Prabhakaran & Rambow, 2013; Rosenthal & McKeown, 2016)

optional:

(Mayfield et al., 2013; Prabhakaran et al., 2013; Jain, 2016; Ding & Pan, 2016; Wei et al., 2016; Chen et al., 2016)

zur Vorbereitung: ((Prabhakaran & Rambow, 2013) oder (Rosenthal & McKeown, 2016)) und ((Le et al., 2011) oder (Mowery et al., 2016))

31.01.2017

Julius Steen

Medizin und NLP – Datenintegration, ...:

(Tsakalidis et al., 2016)

Tamas Janusko

Medizin und NLP – Anonymisierung:

(Uzuner et al., 2007; Szarvas et al., 2007; Wellner & Pustejovsky, 2007; Narayanan & Shmatikov, 2008; Friedlin & McDonald, 2008; Uzuner et al., 2008; Hirschman & Aberdeen, 2010; Benitez & Malin, 2010; Schwartz et al., 2013; Qian & Liu, 2013; Lee et al., 2016; Yadav et al., 2016; Yuwono et al., 2016)

Medizin und NLP – Informationsextraktion, Kommunikation, Vorhersage

(Pitler & Nenkova, 2008; Paul & Dredze, 2011; Ferrucci et al., 2013; Wallace et al., 2013; Chen, 2013; Sarioglu et al., 2013; Paul & Dredze, 2013; Rebholz-Schumann et al., 2013; Teodoro & Naaman, 2013; Navindgi et al., 2016; Jagannatha & Yu, 2016; Althoff et al., 2016; Limsopatham & Collier, 2016; Yancheva & Rudzicz, 2016; Iso et al., 2016; Tsakalidis et al., 2016)

Medizin und NLP – Diagnose, Therapie

(Althoff et al., 2016; Shibata et al., 2016)

Medizin und NLP – Pharmakovigilanz

(Pain et al., 2016; Sarker & Gonzalez, 2016; Huynh et al., 2016)

zur Vorbereitung: Video von Regina Barzilays Keynote bei NAACL'16: "How NLP can help cure cancer?" (<http://techtalks.tv/talks/how-can-nlp-help-cure-cancer/62223/>)

07.02.2017

Diskussion

zur Vorbereitung: Offene Fragen, Kommentare, Kritik, Diskussionsvorschläge zum gesamten Semester (mir bitte zusenden bis spätestens

Optional:

Gesprochene Sprache und Dialogsysteme:

(Shepstone et al., 2013; Johnston et al., 2013; Rakov & Rosenberg, 2013; Traum, 2013; Pérez-Rosas & Mihalcea, 2013; Cummins et al., 2013; Evans et al., 2013; Federico et al., 2013; Kim et al., 2013; Bigot et al., 2013; Hatmi et al., 2013)

Transparenz durch NLP:

(Wilson et al., 2016; Zukov Gregoric et al., 2016; Mihaylov & Nakov, 2016)

Politik – Political Event Extraction:

(Beielser, 2016; Makarov et al., 2016)

Bias:

(Kim et al., 6 91)

Microblogs – Lokalisierung:

(Fink et al., 2009; Gelernter & Mushegian, 2011; Cheng et al., 2013; Schulz et al., 2013; Varga et al., 2013; Han et al., 2013; Crooks et al., 2013; Jurgens, 2013)

Microblogs – Zensur:

(Bamman et al., 2012b; Huang et al., 2013; Xu et al., 2013; Sleeper et al., 2013; Das & Kramer, 2013; Zhu et al., 2013)

Microblogs – Topics:

(Li et al., 2016)

Microblogs – Resolving Nicknames, Entity Morphs:

(Huang et al., 2013; Zhang et al., 2014; 2015; Wang et al., 2016a)

Emotionserkennung:

(Hasegawa et al., 2013; Zhou et al., 2016b; Gui et al., 2016; Summa et al., 2016; ?; Blevins et al., 2016)

Personalisierung:

(Yu et al., 2016; Gong et al., 2016; Zhou et al., 2016c)

Wissenschaftsethik in der Informatik:

(Blue et al., 2014)

Soziale Medien – Uncertainty, Sarkasmus, Ironie, ...

(Vincze, 2016; Joshi et al., 2016; Poria et al., 2016; Zhang et al., 2016a; Van Hee et al.,

2016)

Sentiment Analysis – Foren:

(Qiu et al., 2013; Qiu & Jiang, 2013; Chen et al., 2013; Saeidi et al., 2016; Wang et al., 2016b)

Sentiment Analysis – Politische Meinungen:

(Arunachalam & Sarkar, 2013; Mukherjee & Liu, 2013; Mukherjee et al., 2013; Lin et al., 2013; Bhosale et al., 2013; Cohen & Ruths, 2013; Pavlick et al., 2016; Celli et al., 2016; Menini & Tonelli, 2016; Johnson & Goldwasser, 2016)

Sentiment Analysis – Microblogs:

(Rudra et al., 2016; Teng et al., 2016; Al-Twairesh et al., 2016; Sidarenka & Stede, 2016; Blevins et al., 2016)

Stance Classification – Microblogs:

(Somasundaran & Wiebe, 2009; Bamman et al., 2012a; Augenstein et al., 2016; Ebrahimi et al., 2016; Chen & Ku, 2016; Zubiaga et al., 2016; Johnson & Goldwasser, 2016)

Bemerkungen:

Leistungsnachweise: Lektüre und aktive Teilnahme (1/3), Co-Leitung einer Sitzung (1/3), Hausarbeit (1/3). Hausarbeit: 8-10 Seiten (Proseminar), 12-15 Seiten (Hauptseminar) inkl. Bibliographie. Die Hausarbeit kann auch per Email an mich geschickt werden, aber *nicht* als Word-Datei sondern nur als PDF-Datei. – Ich empfehle, wissenschaftliche Texte mit Latex und Bibtex zu verfassen.

Regelmäßige Teilnahme (d.i. nicht mehr als einmal unentschuldigtes Fehlen) ist Voraussetzung für den Scheinerwerb. Zu jeder Sitzung müssen jeweils zwei Fragen (!) zu einem Papier abgegeben werden, das in der aktuellen Sitzung vorgestellt wird. Abgabe entweder per Email bis spätestens 13 Uhr am Montag vor der Sitzung (oder – im Notfall – schriftlich direkt vor der Sitzung). Dies geht in die Bewertung für aktive Teilnahme am Seminar ein.

Literatur: Viele Papiere können direkt aus der *ACL Anthology* kopiert werden (<http://acl.ldc.upenn.edu/>), insbesondere alle Papiere der (*E/NA*)*ACL*-, *Coling*- und *EMNLP*-Konferenzen, alle Workshops, die im Rahmen dieser Konferenzen veranstaltet wurden und die Zeitschrift *Computational Linguistics*. Papiere, die von der *AAAI* publiziert wurden (*AAAI*-Konferenz, *AAAI*-Workshops, *AAAI*-Symposia, etc.) sind in der *AAAI Digital Library* verfügbar (<http://www.aaai.org/Library>). – Die meisten weiteren Zeitschriften sind elektronisch verfügbar über die UB (<http://rzblx1.uni-regensburg.de/ezeit/search.phtml?bibid=UBHE>) – oder stehen dort im Regal.

Sprechstunde: Auf Vereinbarung (Email, Telefon) bei mir im Büro, ggf. auch im Anschluß an das Seminar.

Hausarbeiten:

Maximal 8-10 Seiten (Proseminar), 12-15 Seiten (Hauptseminar) inkl. Abbildungen, inkl. Literaturverzeichnis.

Inhalt: Fokus auf das vorgestellte Papier; NICHT *Related Work*-Kapitel referieren, wenn die entsprechenden Papiere nicht gelesen wurden; Evaluierung berichten; WICHTIG: mit eigener Meinung oder Bewertung abschließen.

Stil: Wissenschaftlichkeit drückt sich nicht durch lange, komplizierte Sätze und exzessiven Gebrauch von Fremdwörtern aus – deshalb bitte kurze Sätze, einfache Sprache; Hausarbeiten vor der Abgabe Korrektur lesen oder Korrektur lesen lassen (s. auch *Dos and donts: Hinweise zur Abfassung wissenschaftlicher Arbeiten* von Prof. Frank – http://www.cl.uni-heidelberg.de/~frank/materials/dos_and_donts.pdf). Ich schätze Wikipedia als Gegenstand meiner Forschung sehr, nicht aber als Quelle für wissenschaftliche Arbeiten. Hausarbeiten, die Wikipedia (oder auch andere allgemeine Enzyklopädien) als Beleg zitieren, werde ich zurückweisen. Bitte lesen und zitieren Sie Fachliteratur!

Seminararbeit (d.i. eine praktische Arbeit) ist auch möglich. Sollte durch 5-6 Seiten Bericht begleitet werden.

Abgabetermin: bis spätestens 31. März 2017; per Email als PDF-Datei (kein Microsoft Word!) oder ausgedruckt per Post – Matrikelnummer und Studiengang nicht vergessen!

References

- Abu-Jbara, Amjad, Ben King, Mona Diab & Dragomir Radev (2013). Identifying opinion subgroups in Arabic online discussions. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 829–835.
- Agarwal, Apoorv, Anup Kotalwar & Owen Rambow (2013). Automatic extraction of social networks from literary text: A case study on Alice in Wonderland. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, Nagoya, Japan, 14–18 October 2013, pp. 1202–1208.
- Al-Twairsh, Nora, Hend Al-Khalifa & Abdulmalik AlSalman (2016). AraSenTi: Large-scale Twitter-specific Arabic sentiment lexicons. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 697–705.
- Althoff, Tim, Kevin Clark & Jure Leskovec (2016). Large-scale analysis of counseling conversations: An application of natural language processing to mental health. *Transactions of the Association for Computational Linguistics*, 4:463–476.
- Argamon, Shlomo, Mosche Koppel, James W. Pennebaker & Jonathan Schler (2009). Automatically profiling the author of an anonymous text. *Communications of the ACM*, 52(2):119–123.
- Arunachalam, Ravi & Sandipan Sarkar (2013). The new eye of the government: Citizen sentiment analysis in social media. In *Proceedings of the Workshop on Natural Language Processing for Social Media (SocialNLP)*, Nagoya, Japan, 14 October 2013, pp. 23–28.
- Augenstein, Isabelle, Tim Rocktäschel, Andreas Vlachos & Kalina Bontcheva (2016). Stance detection with bidirectional conditional encoding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 876–885.
- Bachenko, Joan, Eileen Fitzpatrick & Michael Schonwetter (2008). Verification and implementation of language-based deception indicators in civil and criminal narratives. In *Proceedings of the 22nd International Conference on Computational Linguistics*, Manchester, U.K., 18–22 August 2008, pp. 41–48.
- Bamman, David, Jacob Eisenstein & Tyler Schnoebelen (2012a). *Gender in Twitter: Styles, stances, and social networks*. Published at arXiv:1210.4567.
- Bamman, David, Brendan O’Connor & Noah A. Smith (2012b). Censorship and content deletion in Chinese social media. *First Monday*, 17(3-5).
- Beieler, John (2016). Generating politically-relevant event data. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 37–42.
- Benitez, Kathleen & Bradley Malin (2010). Evaluating re-identification risks with respect to the HIPAA privacy rule. *Journal of the American Medical Informatics Association*, 17(2):160–177.
- Benson, Edward, Aria Haghighi & Regina Barzilay (2011). Event discovery in social media feeds. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Portland, Oreg., 19–24 June 2011, pp. 389–398.
- Benton, Adrian, Raman Arora & Mark Dredze (2016). Learning multiview embeddings of Twitter users. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 14–19.
- Bergsma, Shane & Benjamin Van Durme (2013). Using conceptual class attributes to characterize social media users. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 710–720.
- Beusmans, Jack & Kären Wieckert (1989). Computing, research, and war: If knowledge is power, where is responsibility? *Communications of the ACM*, 32(8):939–947.
- Bhosale, Shruti, Heath Vinicombe & Raymond Mooney (2013). Detecting promotional content in Wikipedia. In *Proceedings of the 2013 Conference on Empirical Methods in Natural*

- Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1851–1857.
- Bigot, Benjamin, Grégory Senay, Linarès, Corinne Fredouille & Richard Dufour (2013). Combining acoustic name spotting and continuous context models to improve spoken person name recognition in speech. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Blevins, Terra, Robert Kwiatkowski, Jamie MacBeth, Kathleen McKeown, Desmond Patton & Owen Rambow (2016). Automatically processing Tweets from gang-involved youth: Towards detecting loss and aggression. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2196–2206.
- Blodgett, Su Lin, Lisa Green & Brendan O’Connor (2016). Demographic dialectal variation in social media: A case study of African-American English. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1119–1130.
- Blue, Ethan, Michael Levine & Dean Nieuwsma (2014). *Engineering and War: Militarism, Ethics, Institutions, Alternatives*. Morgan and Claypool Publishers.
- Bond, Gary D. & Adrienne Y. Lee (2005). Language of lies in prison: Linguistic classification of prisoners’ truthful and deceptive natural language. *Applied Cognitive Psychology*, 19:313–329.
- Bostrom, Nick und Yudkowsky, Eliezer (2014). The ethics of artificial intelligence. In K. Frankish & W.M. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence*, pp. 316–334. Cambridge, U.K.: Cambridge University Press.
- Bradford, Roger B. (2006). Relationship discovery in large text collections using latent semantic indexing. In *Proceedings of the 4th Workshop on Link Analysis, Counterterrorism, and Security*, Bethesda, Md., 20–22 April 2006.
- Brennan, Michael, Sadia Afroz & Rachel Greenstadt (2012). Adversarial stylometry: Circumventing authorship recognition to preserve privacy and anonymity. *ACM Transactions on Information and System Security*, 15(3). Article No. 12.
- Burger, John D., John Henderson, George Kim & Guido Zarrella (2011). Discriminating gender on Twitter. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, Edinburgh, Scotland, U.K., 27–29 July 2011, pp. 1301–1309.
- Burgoon, Judee K., J.P. Blair, Tiantian Qin & Jr. Nunamaker, Jay F. (2003). Detecting deception through linguistic analysis. In *Proceedings of Intelligence and Security Informatics (ISI)*, p. 958.
- Burri, Thomas (2016). Machine learning and the law: Five theses. In *Proceedings of the NIPS 2016 Workshop on Machine Learning and the Law*, Barcelona, Spain, 8 December 2016.
- Cano, Elizabeth, Yuland He, Kang Liu & Jun Zhao (2013). A weakly supervised Bayesian model for violence detection in social media. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, Nagoya, Japan, 14–18 October 2013, pp. 109–117.
- Celli, Fabio, Evgeny Stepanov, Massimo Poesio & Giuseppe Riccardi (2016). Predicting Brexit: Classifying agreement is better than sentiment and pollsters. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016, pp. 110–118.
- Chen, Annie (2013). Patient experience in online support forums: Modeling interpersonal interactions and medication use. In *51st Annual Meeting of the Association for Computational Linguistics: Proceedings of the Student Research Workshop*, Sofia, Bulgaria, 5–7 August 2013, pp. 16–22.
- Chen, Chengyao, Zhitao Wang, Yu Lei & Wenjie Li (2016). Content-based influence modeling for opinion behavior prediction. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2207–2216.
- Chen, Wei-Fan & Lun-Wei Ku (2016). UTCNN: A deep learning model of stance classification on social media text. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 1635–1645.
- Chen, Zhiyuan, Bing Liu, Meichun Hsu, Malu Castellanos & Riddhiman Ghosh (2013). Iden-

- tifying intention posts in discussion forums. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 1041–1050.
- Cheng, Zhiyuan, James Caverlee & Kyumin Lee (2013). A content-driven framework for geolocating microblog users. *ACM Transactions on Intelligent Systems and Technology*, 4(1). Article 2 (27 pages).
- Cho, Adrian (2013). Network science at center of surveillance dispute. *Science*, 340(6138):1272.
- Chua, Freddy Chong Tat & Sitaram Asur (2013). Automatic summarization of events from social media. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Ciot, Morgane, Morgan Sonderegger & Derek Ruths (2013). Gender inference of Twitter users in non-English contexts. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1136–1145.
- Coffman, Thayne, Seth Greenblatt & Sherry Marcus (2004). Graph-based technologies for intelligence analysis. *Communications of the ACM*, 47(3):45–47.
- Cohen, Raviv & Derek Ruths (2013). Classifying political orientation on Twitter: It’s not easy! In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Cohn, Michael A., Matthias R. Mehl & James W. Pennebaker (2004). Linguistic markers of psychological change surrounding September 11, 2001. *Psychological Science*, 15(10):687–693.
- Coppersmith, Glen, Kristy Hollingshead, H. Andrew Schwartz, Molly Ireland, Rebecca Resnik, Kate Loveys, April Foreman & Loring Ingraham (2016). The clinical panel: Leveraging psychological expertise during NLP research. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 132–137.
- Crooks, Andrew, Arie Croitoru, Anthony Stefanidis & Jacek Radzikowski (2013). #earthquake: Twitter as a distributed sensor system. *Transactions in GIS*, 17(1):124–147.
- Cummins, Nicholas, Julien Epps, Vidhyasaharan Sethu, Michael Breakspear & Roland Goecke (2013). Modeling spectral variability for the classification of depressed speech. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Das, Sauvik & Adam Kramer (2013). Self-censorship on Facebook. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- De Choudhury, Munmun, Michael Gamon, Scott Counts & Eric Horvitz (2013). Predicting depression via social media. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Diao, Qiming & Jing Jiang (2013). A unified model for topics, events and users on Twitter. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1869–1879.
- Ding, Tao & Shimei Pan (2016). Personalized emphasis framing for persuasive message generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1432–1441.
- Dredze, Mark, Nicholas Andrews & Jay DeYoung (2016). Twitter at the Grammys: A social media corpus for entity linking and disambiguation. In *Proceedings of the 4th Workshop on Natural Language Processing and Social Media*, Austin, Texas, 1 November 2016, pp. 20–25.
- Ebrahimi, Javid, Dejing Dou & Daniel Lowd (2016). Weakly supervised Tweet stance classification by relational bootstrapping. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1012–1017.
- Edsall, John T. (1975). Scientific freedom and responsibility. report of the AAAS committee on scientific freedom and responsibility. *Science*, 188(4189):687–693.
- Eisenstein, Jacob (2013). What to do about bad language on the internet. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 359–369.

- Eisenstein, Jacob, Noah A. Smith & Eric P. Xing (2011). Discovering sociolinguistic associations with structured sparsity. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Portland, Oreg., 19–24 June 2011, pp. 1365–1374.
- El-Arini, Khalid, Min Xu, Emily Fox & Carlos Guestrin (2013). Representing documents through their readers. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Chicago, Ill., 11–14 August 2013, pp. 14–22.
- Elson, David K., Nicholas Dames & Kathleen R. McKeown (2010). Extracting social networks from literary fiction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Uppsala, Sweden, 11–16 July 2010, pp. 138–147.
- Evans, Nicholas, Tomi Kinnunen & Junichi Yamagishi (2013). Spoofing and countermeasures for automatic speaker verification. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Fast, Ethan & Eric Horvitz (2016). Identifying dogmatism in social media: Signals and models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 690–699.
- Federico, Alegre, Ravichander Vippera, Asmaa Amehraye & Nicholas Evans (2013). A new speaker verification spoofing countermeasure based on local binary patterns. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Feng, Vanessa Wei & Graeme Hirst (2013). Detecting deceptive opinions with profile compatibility. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, Nagoya, Japan, 14–18 October 2013, pp. 338–346.
- Ferrucci, David A., Anthony Levas, Sugato Bagchi, David Gondek & Erik T. Mueller (2013). Watson: Beyond Jeopardy! *Artificial Intelligence*, 199-200:93–105.
- Fink, Clay, Christine Piatko, James Mayfield, Tim Finin & Justin Martineau (2009). Geolocating blogs from their textual content. In *Working Notes of the AAAI Spring Symposium on Social Semantic Web: Where Web 2.0 Meets Web 3.0*, Stanford, Cal., 23–25 March 2009.
- Flekova, Lucie, Jordan Carpenter, Salvatore Giorgi, Lyle Ungar & Daniel Preoțiu-Pietro (2016a). Analyzing biases in human perception of user age and gender from text. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016.
- Flekova, Lucie & Iryna Gurevych (2013). Can we hide in the web? large scale simultaneous age and gender author profiling in social media – Notebook for PAN at CLEF 2013. In *Proceedings of CLEF 2013 Labs and Workshops – Notebook Papers*, Valencia, Spain, 23-26 September 2013.
- Flekova, Lucie, Daniel Preoțiu-Pietro & Lyle Ungar (2016b). Exploring stylistic variation with age and income on Twitter. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 313–319.
- Friedlin, F. Jeff & Clement J. McDonald (2008). A software tool for removing patient identifying information from clinical documents. *Journal of the American Medical Informatics Association*, 15(5):601–610.
- Gelernter, Judith & Nikolai Mushegian (2011). Geo-parsing messages from microtext. *Transactions in GIS*, 15(6):753–773.
- Glantz, Stanton A. & Norm V. Albers (1974). Department of Defense R & D in the university. *Science*, 186(4165):706–711.
- Gong, Lin, Mohammad Al Boni & Hongning Wang (2016). Modeling social norms evolution for personalized sentiment classification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 855–865.
- Graciarena, Martin, Elizabeth Shriberg, Andreas Stolcke, Frank Enos, Julia Hirschberg & Sachin Kajarekar (2006). Combining prosodic, lexical and cepstral systems for deceptive speech detection. In *Proceedings of the 2006 IEEE International Conference on Acoustics, Speech,*

- and *Signal Processing*, Toulouse, France, 15–19 June 2006, pp. 1033–1036.
- Greenwald, Glenn (2014). *No Place to Hide: Edward Snowden, the NSA and the U.S. Surveillance State*. New York, N.Y.: Metropolitan Books.
- Grinberg, Nir, Mor Naaman, Blake Shaw & Gilad Lotan (2013). Extracting diurnal patterns of real world activity from social media. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Gui, Lin, Dongyin Wu, Ruifeng Xu, Qin Lu & Yu Zhou (2016). Event-driven emotion cause extraction with corpus construction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1639–1649.
- Guo, Weiwei, Hao Li, Heng Ji & Mona Diab (2013a). Linking tweets to news: A framework to enrich short text data in social media. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 239–249.
- Guo, Yuhang, Bing Qin, Ting Liu & Sheng Li (2013b). Microblog entity linking by leveraging extra posts. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 863–868.
- Hajaj, Chen, Noam Hazon, David Sarne & Avshalom Elmalech (2013). Search more, disclose less. In *Proceedings of the 27th Conference on the Advancement of Artificial Intelligence*, Bellevue, Wash., 14–18 July 2013, p. ??
- Han, Bo, Paul Cook & Timothy Baldwin (2013). A stacking-based approach to Twitter user geolocation prediction. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, Sofia, Bulgaria, 4–9 August 2013, pp. 7–12.
- Hancock, Jeffrey T., Lauren Curry, Saurabh Goorha & Michael Woodworth (2005). Automated linguistic analysis of deceptive and truthful synchronous computer-mediated communication. In *Proceedings of the Hawaii International Conference on System Sciences (HICSS)*, p. 22c.
- Hasegawa, Takayuki, Nobuhiro Kaji, Naoki Yoshinaga & Masashi Toyoda (2013). Predicting and eliciting addressee’s emotion in online dialogue. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 964–972.
- Hatmi, Mohamd, Christine Jacquin, Emmanuel Morin & Sylvain Meignier (2013). Incorporating named entity recognition into the speech transcription process. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Hirschman, Lynette & John Aberdeen (2010). Measuring risk and information preservation: Toward new metrics for de-identification of clinical texts. In *Proceedings of the 2nd Louhi Workshop on Text and Data Mining of Health Documents*, Los Angeles, Cal., 5 June 2010, pp. 72–75.
- Holden, Constance (1975). Privacy: Congressional efforts are coming to fruition. *Science*, 188(4189):713–715.
- Hovy, Dirk & Shannon L. Spruit (2016). The social impact of Natural Language Processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 591–598.
- Huang, Hongzhao, Zhen Wen, Dian Yu, Heng Ji, Yizhou Sun, Jiawei Han & He Li (2013). Resolving entity morphs in censored data. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 1083–1093.
- Huynh, Trung, Yulan He, Alistair Willis & Stefan R ger (2016). Adverse drug reaction classification with deep neural networks. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 877–887.
- Iso, Hayate, Shoko Wakamiya & Eiji Aramaki (2016). Forecasting word model: Twitter-based influenza surveillance and prediction. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 76–86.
- Jagannatha, Abhyuday & Hong Yu (2016). Structured prediction models for RNN based se-

- quence labeling in clinical text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 856–865.
- Jain, Siddharth (2016). Identifying sensible participants in online discussions. In *Proceedings of the 4th Workshop on Natural Language Processing and Social Media*, Austin, Texas, 1 November 2016, pp. 41–47.
- Ji, Heng & Ralph Grishman (2011). Knowledge base population: Successful approaches and challenges. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Portland, Oreg., 19–24 June 2011, pp. 1148–1158.
- Ji, Heng, Joel Nothman, Ben Hachey & Radu Florian (2016). Overview of TAC-KBP2015 trilingual entity discovery and linking. In *Proceedings of the Text Analysis Conference*, National Institute of Standards and Technology, Gaithersburg, Maryland, USA, 16–17 November 2015.
- Johnson, Kristen & Dan Goldwasser (2016). "All I know about politics is what I read in Twitter": Weakly supervised models for extracting politicians' stances from Twitter. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2966–2977.
- Johnston, Michael, Patrick Ehlen, Frederick G. Conrad, Michael F. Schober, Christopher Antoun, Stefanie Fail, Andrew Hupp, Lucas Vickers, Huiying Yan & Chan Zhang (2013). Spoken dialog systems for automated survey interviewing. In *Proceedings of the SIGdial 2013 Conference: The 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Metz, France, 22–24 August 2013, pp. 329–333.
- Joshi, Aditya, Vaibhav Tripathi, Kevin Patel, Pushpak Bhattacharyya & Mark Carman (2016). Are word embedding-based features useful for sarcasm detection? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1006–1011.
- Jurgens, David (2013). That's what friends are for: Inferring location in online social media platforms based on social relationships. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Kairam, Sanjay Ram, Meredith Ringel Morris, Jaime Teevan, Dan Liebling & Susan Dumais (2013). Towards supporting search over trending events with social media. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Kamijo, Koichi, Tetsuya Nasukawa & Hideya Kitamura (2016). Personality estimation from Japanese text. In *Proceedings of the Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016, pp. 101–109.
- Kim, Samuel, Fabio Valente & Alessandro Vinciarelli (2013). Annotation and detection of conflict escalation in political debates. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Kim, Sunghwan Mac, Stephen Wan & Cecile Paris (2016). Detecting social roles in Twitter. In *Proceedings of the 4th Workshop on Natural Language Processing and Social Media*, Austin, Texas, 1 November 2016, pp. 34–40.
- Kim, Sunghwan Mac, Stephen Wan, Cecile Paris, Jin Brian & Bella Robinson (86-91). The effects of data collection methods in Twitter. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016.
- Kramer, Adam D. I., Jamie E. Guillory & Jeffrey T. Hancock (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790.
- Kramer, Adam D.I. (2012). The spread of emotion via Facebook. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, Austin, Texas, 5-10 May 2012, pp. 767–770.
- Lamb, Alex, Michael J. Paul & Mark Dredze (2013). Separating fact from fear: Tracking flu infections on Twitter. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 789–795.

- Le, Xuan, Ian Lancashire, Graeme Hirst & Regina Jokel (2011). Longitudinal detection of dementia through lexical and syntactic changes in writing: A case study of three British novelists. *Literary and Linguistic Computing*, 26(4):435–461.
- Lee, Ji Young, Franck Dernoncourt, Ozlem Uzuner & Peter Szolovits (2016). Feature-augmented neural networks for patient note de-identification. In *Proceedings of the Clinical Natural Language Processing Workshop*, Osaka, Japan, 11 December 2016, pp. 17–22.
- Li, Chiwei, Myle Ott & Claire Cardie (2013a). Identifying manipulated offerings on review portals. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1933–1942.
- Li, Fangtao, Yang Gao, Shuchang Zhou, Xiance Si & Decheng Dai (2013b). Deceptive answer prediction with user preference graph. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 1723–1732.
- Li, Jing, Ming Liao, Wei Gao, Yulan He & Kam-Fai Wong (2016). Topic extraction from microblog posts using conversation structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 2114–2123.
- Li, Jiwei, Claire Cardie & Sujian Li (2013c). TopicSpam: A topic-model based approach for spam detection. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 217–221.
- Limsopatham, Nut & Nigel Collier (2016). Normalising medical concepts in social media texts by learning semantic representation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 1014–1023.
- Lin, Ching-Sheng, Samira Shaikh, Jennifer Stromer-Galley, Jennifer Crowley, Tomek Strzalkowski & Veena Ravishankar (2013). Topical positioning: A new method for predicting opinion changes in conversation. In *Proceedings of the Workshop on Language Analysis in Social Media*, Atlanta, Georgia, 13 June 2013, pp. 41–48.
- Litvak, Marina, Jahna Otterbacher, Chee Siang Ang & David Atkins (2016). Social and linguistic behavior and its correlation to trait empathy. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016.
- Liu, Fei, Julien Perez & Scott Nowson (2016a). A recurrent and compositional model for personality trait recognition from short texts. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016, pp. 20–29.
- Liu, Tong, Christopher Homan, Cecilia Ovesdotter Alm, Megan Lytle, Ann Marie White & Henry Kautz (2016b). Understanding discourse on work and job-related well-being in public social media. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 1044–1053.
- Liu, Xiaohua, Yitong Li, Haocheng Wu, Ming Zhou, Furu Wei & Yi Lu (2013). Entity linking for tweets. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 1304–1311.
- Lott, P.R., S. Guggenbühl, A. Schneeberger, A.E. Pulver & H.H. Stassen (2002). Linguistic analysis of the speech output of schizophrenic, bipolar, and depressive patients. *Psychopathology*, 35:220–227.
- Makarov, Peter, Jasmine Lorenzini & Hanspeter Kriesi (2016). Constructing an annotated corpus for protest event mining. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 102–107.
- Mayfield, Elijah, David Adamson & Carolyn Penstein Rosé (2013). Recognizing rare social phenomena in conversation: Empowerment detection in support group chatrooms. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 104–113.
- Mehdad, Yashar & Joel Tetreault (2016). Do characters abuse more than words? In *Proceedings*

- of the SIGdial 2016 Conference: The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, Los Angeles, Calif., 13–15 September 2016, pp. 299–303.
- Menini, Stefano & Sara Tonelli (2016). Agreement and disagreement: Comparison of points of view in the political domain. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2461–2470.
- Mihaylov, Todor & Preslav Nakov (2016). Hunting for troll comments in news community forums. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 399–405.
- Mowery, Danielle L., Albert Park, Craig Bryan & Mike Conway (2016). Towards automatically classifying depressive symptoms from Twitter data for population health. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016, pp. 182–191.
- Mukherjee, Arjun & Bing Liu (2013). Discovering user interactions in ideological discussions. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 671–681.
- Mukherjee, Arjun, Vivek Venkataraman, Bing Liu & Sharon Meraz (2013). Public dialogue: Analysis of tolerance in online discussions. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 1680–1690.
- Narayanan, Arvind & Vitaly Shmatikov (2008). Robust de-anonymization of large sparse datasets. In *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, Cal., 18–21 May 2008, pp. 111–125.
- Navindgi, Amit, Caroline Brun, Cécile Boulard Masson & Scott Nowson (2016). Steps toward automatic understanding of the function of affective language in support groups. In *Proceedings of the 4th Workshop on Natural Language Processing and Social Media*, Austin, Texas, 1 November 2016, pp. 26–33.
- Nguyen, Dong, Rilana Gravel, Dolf Trieschnigg & Theo Meder (2013a). “how old do you think I am?” A study of language and age in Twitter. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Nguyen, Thin, Bo Dao, Dinh Phung, Svetha Venkatesh & Michael Berk (2013b). Online social capital: Mood, topical and psycholinguistic aspects. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Nitta, Taisei, Fumito Masui, Michal Ptaszynski, Yasutomo Kimura, Rafal Rzepka & Kenji Araki (2013). Detecting cyberbullying entries on informal school websites based on category relevance maximization. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, Nagoya, Japan, 14–18 October 2013, pp. 579–586.
- Ott, Myle, Claire Cardie & Jeffrey T. Hancock (2013). Negative deceptive opinion spam. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 497–501.
- Ott, Myle, Yejin Choi, Claire Cardie & Jeffrey T. Hancock (2011). Finding deceptive opinion spam by any stretch of imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Portland, Oreg., 19–24 June 2011, pp. 309–319.
- Pain, Julie, Jessie Levacher, Adam Quinquenel & Anja Belz (2016). Analysis of Twitter data for postmarketing surveillance in pharmacovigilance. In *Proceedings of the 2nd Workshop on Noisy User-generated Text*, Osaka, Japan, 11 December 2016, pp. 94–101.
- Paul, Michael J. & Mark Dredze (2011). You are what you tweet: Analyzing Twitter for public health. In *Proceedings of the 5th International Conference on Weblogs and Social Media*, Barcelona, Spain, 17–21 July 2011.
- Paul, Michael J. & Mark Dredze (2013). Drug extraction from the web: Summarizing drug experiences with multi-dimensional topic models. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 168–178.

- Pavlick, Ellie, Heng Ji, Xiaoman Pan & Chris Callison-Burch (2016). The gun violence database: A new task and data set for NLP. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1018–1024.
- Peng, Nanyun & Mark Dredze (2016). Improving named entity recognition for Chinese social media with word segmentation representation learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 149–155.
- Pérez-Rosas, Verónica & Rada Mihalcea (2013). Sentiment analysis of online spoken reviews. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Pestian, John P., Pawel Matykiewicz, Michelle Linn-Gust, Brett South, Ozlem Uzuner, Jan Wiebe, K. Bretonnel Cohen, John Hurdle & Christopher Brew (2012). Sentiment analysis of suicide notes: A shared task. *Biomedical Informatics Insights*, 5(Suppl.1):3–16.
- Pitler, Emily & Ani Nenkova (2008). Revisiting readability: A unified framework for predicting text quality. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Waikiki, Honolulu, Hawaii, 25–27 October 2008, pp. 186–195.
- Popp, Robert, Thomas Armour, Ted Senator & Kristen Numrzech (2004). Countering terrorism through information technology. *Communications of the ACM*, 47(3):36–43.
- Poria, Soujanya, Erik Cambria, Devamanyu Hazarika & Prateek Vij (2016). A deeper look into sarcastic Tweets using deep convolutional neural networks. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 1601–1612.
- Poulston, Adam, Mark Stevenson & Kalina Bontcheva (2016). User profiling with geo-located posts and demographic data. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 43–48.
- Prabhakaran, Vinodkumar, Ajita John & Dorée D. Seligman (2013). Who had the upper hand? Ranking participants of interactions based on their relative power. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, Nagoya, Japan, 14–18 October 2013, pp. 365–373.
- Prabhakaran, Vinodkumar & Owen Rambow (2013). Written dialog and social power: Manifestations of different types of power in dialog behavior. In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, Nagoya, Japan, 14–18 October 2013, pp. 216–224.
- Ptaszynski, Michal, Fumito Masui, Taisei Nitta, Suzuha Hatakeyama, Yasutomo Kimura, Rafal Rzepka & Kenji Araki (2016). Sustainable cyberbullying detection with category-maximized relevance of harmful phrases and double-filtered automatic optimization. *International Journal of Child-Computer Interaction*, 8:15–30.
- Qian, Tiejun & Bing Liu (2013). Identifying multiple userids of the same author. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1124–1135.
- Qiu, Minghui & Jing Jiang (2013). A latent variable model for viewpoint discovery from threaded forum posts. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 1031–1040.
- Qiu, Minghui, Liu Yang & Jing Jiang (2013). Mining user relations from online discussions using sentiment analysis and probabilistic matrix factorization. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 401–410.
- Rakov, Rachel & Andrew Rosenberg (2013). "Sure I did the right thing": A system for sarcasm detection in speech. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Rangel, Francisco, Paolo Rosso, Moshe Koppel, Efstathios Stamatatos & Giacomo Inches (2013). Overview of the author profiling task at PAN 2013. In *Proceedings of CLEF 2013 Labs and Workshops – Notebook Papers*, Valencia, Spain, 23–26 September 2013.

- Rao, Delip, Michael Paul, Clayton Fink, David Yarowsky, Timothy Oates & Glen Coppersmith (2011). Hierarchical Bayesian models for latent attribute detection in social networks. In *Proceedings of the 5th International Conference on Weblogs and Social Media*, Barcelona, Spain, 17–21 July 2011.
- Rebholz-Schumann, Dietrich, Simon Clematide, Fabio Rinaldi, Senay Kafkas, Erik M. van Muligen, Chinh Bui, Johannes Hellrich, Ian Lewin, David Milward, Michael Poprat, Antonio Jimeno-Yepes, Udo Hahn & Jan A. Kors (2013). Multilingual semantic resources and parallel corpora in the biomedical domain: The CLEF-ER challenge. In *Proceedings of CLEF 2013 Labs and Workshops – Notebook Papers*, Valencia, Spain, 23–26 September 2013.
- Reddy, Sravana & Kevin Knight (2016). Obfuscating gender in social media writing. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 17–26.
- Resnik, Philip, Anderson Garron & Rebecca Resnik (2013). Using topic modeling to improve prediction of neuroticism and depression in college students. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1348–1353.
- Rocha, Andersona, Walter Scheirer, Christopher Forstall, Thiago Cavalcante, Antonio Theophilo, Bingyu Shen, Ariadne Carvalho & Efstathios Stamatatos (2017). Authorship attribution for social media forensics. *IEEE Transactions on Information Forensics and Security*, ??(??):??
- Rosenthal, Sara & Kathleen McKeown (2016). Social proof: The impact of author traits on influence detection. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 27–36.
- Rubenstein, Ira S., Ronald D. Lee & Paul M. Schwartz (2008). Data mining and internet profiling: Emerging regulatory and technological approaches. *The University of Chicago Law Review*, 75(1):261–285.
- Rude, Stephanie S., Eva-Maria Gortner & James W. Pennebaker (2004). Language use of depressed and depression-vulnerable college students. *Cognition and Emotion*, 18(8):1121–1133.
- Rudra, Koustav, Shruti Rijhwani, Rafiya Begum, Kalika Bali, Monojit Choudhury & Niloy Ganguly (2016). Understanding language preference for expression of opinion and sentiment: What do Hindi-English speakers do on Twitter? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1131–1141.
- Sadeque, Farig, Ted Pedersen, Thamar Solorio, Prasha Shrestha, Nicolas Rey-Villamizar & Steven Bethard (2016). Why do they leave: Modeling participation in online depression forums. In *Proceedings of the 4th Workshop on Natural Language Processing and Social Media*, Austin, Texas, 1 November 2016, pp. 14–19.
- Saeidi, Marzieh, Guillaume Bouchard, Maria Liakata & Sebastian Riedel (2016). SentiHood: Targeted aspect based sentiment analysis dataset for urban neighbourhoods. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 1546–1556.
- Sarioglu, Efsun, Kabir Yadav & Hyeong-Ah Choi (2013). Topic modeling based classification of clinical reports. In *51st Annual Meeting of the Association for Computational Linguistics: Proceedings of the Student Research Workshop*, Sofia, Bulgaria, 5–7 August 2013, pp. 67–73.
- Sarker, Abeed & Graciela Gonzalez (2016). Data, tools and resources for mining social media drug chatter. In *Proceedings of the Fifth Workshop on Building and Evaluating Resources for Biomedical Text Mining*, Osaka, Japan, 12 December 2016, pp. 99–107.
- Schmidt, Eric & Jared Cohen (2013). *The New Digital Age: Reshaping the Future of People, Nations and Business*. New York, N.Y.: Alfred A. Knopf.
- Schneier, Bruce (2014). Metadata = surveillance. *IEEE Security and Privacy*, 12(2):84.
- Schneier, Bruce (2015). *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World*. New York, N.Y.: W.W. Norton and Company.
- Schuler, Douglas & Jonathan Jacky (1989). Responsibility. *Communcation of the ACM*,

32(8):925–927.

- Schulz, Axel, Aristotelis Hadjakos, Heiko Paulheim, Johannes Nachtwey & Max Mühlhäuser (2013). A multi-indicator approach for geolocalization of tweets. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Schwartz, Roy, Oren Tsr, Ari Rappoport & Moshe Koppel (2013). Authorship attribution of micro-messages. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1880–1891.
- Shapley, Deborah (1972). Defense research: The names are changed to protect the innocent. *Science*, 175(4024):866–868.
- Shepstone, Sven Ewan, Zheng-Hua Tan & Søren Holdt Jensen (2013). Demographic recommendation by means of group profile elicitation using speaker age and gender recognition. In *Proceedings of the 14th Annual Conference of the International Speech Communication Association*, Lyon, France, 25–29 August 2013.
- Shibata, Daisaku, Shoko Wakamiya, Ayae Kinoshita & Eiji Aramaki (2016). Detecting Japanese patients with Alzheimer’s disease based on word category frequencies. In *Proceedings of the Clinical Natural Language Processing Workshop*, Osaka, Japan, 11 December 2016, pp. 78–85.
- Sidarenka, Uladzimir & Manfred Stede (2016). Generating sentiment lexicons for German Twitter. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016, pp. 80–90.
- Sleeper, Many, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese & Lorie Faith Cranor (2013). The post that wasn’t: Exploring self-censorship on Facebook. In *Proceedings of the Conference on Computer Supported Cooperative Work*, San Antonio, Tex., 23–27 February 2013, pp. 793–802.
- Smith, Laura, Salvatore Giorgi, Rishi Solanki, Johannes Eichstaedt, H. Andrew Schwartz, Muhammad Abdul-Mageed, Anneke Buffone & Lyle Ungar (2016). Does ‘well-being’ translate on twitter? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 2042–2047.
- Somasundaran, Swapna & Janyce Wiebe (2009). Recognizing stances in online debates. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing*, Singapore, 2–7 August 2009.
- Sousa Silva, Rui, Gustavo Laboreiro, Luís Sarmiento, Tim Grant, Eugénio Oliveira & Belinda Maia (2011). ¡@twazn me!!!;¡ Automatic authorship analysis of micro-blogging messages. In *Proceedings of the Conference on Natural Language and Database Systems*, Alicante, Spain, 28–30 June 2011, pp. 161–168.
- Stirman, Shannon Wiltsey & James W. Pennebaker (2001). Word use in the poetry of suicidal and nonsuicidal poets. *Psychosomatic Medicine*, 63:517–522.
- Strauss, Benjamin, Bethany Toma, Alan Ritter, Marie-Catherine de Marneffe & Wei Xu (2016). Results of the WNUT16 named entity recognition shared task. In *Proceedings of the 2nd Workshop on Noisy User-generated Text*, Osaka, Japan, 11 December 2016, pp. 138–144.
- Summa, Anja, Bernd Resch & Michael Strube (2016). Microblog emotion classification by computing similarity in text, time, and space. In *Proceedings of the Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Osaka, Japan, 12 December 2016, pp. 153–162.
- Szarvas, György, Richárd Farkas & Róbert Busa-Fekete (2007). State-of-the-art anonymization of medical records using an iterative machine learning framework. *Journal of the American Medical Informatics Association*, 14(5):574–580.
- Takase, Sho, Akiko Murakami, Miki Enoki, Naoaki Okazaki & Kentaro Inui (2013). Detecting chronic critics based on sentiment polarity and user’s behavior in social media. In *51st Annual Meeting of the Association for Computational Linguistics: Proceedings of the Student Research Workshop*, Sofia, Bulgaria, 5–7 August 2013, pp. 110–116.
- Teng, Zhiyang, Duy Tin Vo & Yue Zhang (2016). Context-sensitive lexicon features for neural sentiment analysis. In *Proceedings of the 2016 Conference on Empirical Methods in Natural*

- Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1629–1638.
- Teodoro, Rannie & Mor Naaman (2013). Fitter with Twitter: Understanding personal health and fitness activity in social media. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Thompson, Clark (1986). Military direction of academic CS research. *Communications of the ACM*, 29(7):583–585.
- Traum, David (2013). Non-cooperative and deceptive virtual agents. *IEEE Intelligent Systems*, 27(6):66–69.
- Tsakalidis, Adam, Maria Liakata, Theo Damoulas, Brigitte Jellinek, Weisi Guo & Alexandra Cristea (2016). Combining heterogeneous user generated data to sense well-being. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 3007–3018.
- Tsur, Oren & Ari Rappoport (2013). Efficient clustering of short messages into general domains. In *Proceedings of the 7th International Conference on Weblogs and Social Media*, Cambridge, Mass., 8–11 July 2013.
- Uzuner, Özlem, Yuan Luo & Peter Szolovits (2007). Evaluating the state-of-the-art in automatic de-identification. *Journal of the American Medical Informatics Association*, 14(5):550–563.
- Uzuner, Özlem, Tawanda C. Sibanda, Yuan Luo & Peter Szolovits (2008). A de-identifier for medical discharge summaries. *Artificial Intelligence in Medicine*, 42(1):13–35.
- Van Hee, Cynthia, Els Lefever & Veronique Hoste (2016). Monday mornings are my fave :) #not exploring the automatic recognition of irony in English tweets. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2730–2739.
- Varga, István, Motoki Sano, Kentaro Torisawa, Chikara Hashimoto, Kiyonori Ohtake, Takao Kawai, Jong-Hoon Oh & Stijn De Saeger (2013). Aid is out there: Looking for help from tweets during a large scale disaster. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, 4–9 August 2013, pp. 1619–1629.
- Verma, Inder M. (2014). Editorial expression of concern: Experimental evidence of massive scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(29):10779.
- Vincze, Veronika (2016). Detecting uncertainty cues in Hungarian social media texts. In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*, Osaka, Japan, 12 December 2016, pp. 11–21.
- Volkova, Svitlana & Yoram Bachrach (2016). Inferring perceived demographics from user emotional tone and user-environment emotional contrast. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 1567–1578.
- Wallace, Byron C., Thomas A Trikalinos, M. Barton Laws, Ira B. Wilson & Eugene Charniak (2013). A generative joint, additive, sequential model of topics and speech acts in patient-doctor communication. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 1765–1775.
- Wang, Wei-Chung, Hung-Chen Chen, Zhi-Kai Ji, Hui-I Hsiao, Yu-Shian Chiu & Lun-Wei Ku (2016a). Whose nickname is this? recognizing politicians from their aliases. In *Proceedings of the 2nd Workshop on Noisy User-generated Text*, Osaka, Japan, 11 December 2016, pp. 61–69.
- Wang, Xingyou, Weijie Jiang & Zhiyong Luo (2016b). Combination of convolutional and recurrent neural network for sentiment analysis of short texts. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2428–2437.
- Wang, Zhongqing, Shoushan Li, Fang Kong & Guodong Zhou (2013). Collective personal profile summarization with social networks. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pp. 715–725.
- Warner, William & Julia Hirschberg (2012). Detecting hate speech on the World Wide Web. In

- Proceedings of the 2nd Workshop on Language in Social Media*, Montréal, Québec, Canada, 7 June 2012, pp. 19–26.
- Waseem, Zeerak (2016). Are you a racist or am I seeing things? Annotator influence on hate speech detection on Twitter. In *Proceedings of the 1st Workshop on Natural Language Processing and Computational Social Science*, Austin, Texas, 5 November 2016, pp. 138–142.
- Waseem, Zeerak & Dirk Hovy (2016). Hateful symbols or hateful people? Predictive features for hate speech detection on Twitter. In *Proceedings of the NAACL Student Research Workshop*, San Diego, Cal., 12–17 June 2016, pp. 88–93.
- Wei, Zhongyu, Yang Liu & Yi Li (2016). Is this post persuasive? Ranking argumentative comments in online forum. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 195–200.
- Wellner, Ben & James Pustejovsky (2007). Automatically identifying the arguments of discourse connectives. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Language Learning*, Prague, Czech Republic, 28–30 June 2007, pp. 92–101.
- Wilson, Shomir, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Chervirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N. Cameron Russell, Thomas B. Norton, Eduard Hovy, Joel Reidenberg & Norman Sadeh (2016). The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 1330–1340.
- Winograd, Terry (1991). Strategic computing research and the universities. In C. Dunlop & R. Kling (Eds.), *Computerization and Controversy: Value Conflicts and Social Choices*, pp. 704–716. San Diego, Cal.: Academic Press Professional.
- Wohlgenannt, Gerhard, Ekaterina Chernyak & Dmitry Ilvovsky (2016). Extracting social networks from literary text with word embedding tools. In *Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities*, Osaka, Japan, 11 December 2016, pp. 18–25.
- Xiao, Yang, Yuan Wang, Hangyu Mao & Zhen Xiao (2016). Predicting restaurant consumption level through social media footprints. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 3328–3338.
- Xu, Jun-Ming, Benjamin Burchfiel, Xiaojin Zhu & Amy Bellmore (2013). An examination of regret in bullying tweets. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, 9–14 June 2013, pp. 697–702.
- Yadav, Shweta, Asif Ekbal, Sriparna Saha & Pushpak Bhattacharyya (2016). Deep learning architecture for patient data de-identification in clinical records. In *Proceedings of the Clinical Natural Language Processing Workshop*, Osaka, Japan, 11 December 2016, pp. 32–41.
- Yancheva, Maria & Frank Rudzicz (2016). Vector-space topic models for detecting Alzheimer’s disease. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 2337–2346.
- Yang, Yi, Ming-Wei Chang & Jacob Eisenstein (2016). Toward socially-infused information extraction: Embedding authors, mentions, and entities. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 1452–1461.
- Yen, John (2004). Emerging technologies for homeland security. *Communications of the ACM*, 47(3):33–35.
- Yu, Yang, Xiaojun Wan & Xinjie Zhou (2016). User embedding for scholarly microblog recommendation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 449–453.
- Yuwono, Steven Kester, Hwee Tou Ng & Kee Yuan Ngiam (2016). Automated anonymization as spelling variant detection. In *Proceedings of the Clinical Natural Language Processing Workshop*, Osaka, Japan, 11 December 2016, pp. 99–103.

- Zhang, Boliang, Hongzhao Huang, Xiaoman Pan, Heng Ji, Kevin Knight, Zhen Wen, Yizhou Sun, Jiawei Han & Bulent Yener (2014). Be appropriate and funny: Automatic entity morph encoding. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Baltimore, Md., 22–27 June 2014, pp. 706–711.
- Zhang, Boliang, Hongzhao Huang, Xiaoman Pan, Sujian Li, Chin-Yew Lin, Heng Ji, Kevin Knight, Zhen Wen, Yizhou Sun, Jiawei Han & Bulent Yener (2015). Context-aware entity morph decoding. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Beijing, China, 26–31 July 2015, pp. 586–595.
- Zhang, Meishan, Yue Zhang & Guohong Fu (2016a). Tweet sarcasm detection using deep neural network. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2449–2460.
- Zhang, Qi, Yang Wang, Yeyun Gong & Xuanjing Huang (2016b). Keyphrase extraction using deep recurrent neural networks on Twitter. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 836–845.
- Zhou, Deyu, Tianmeng Gao & Yulan He (2016a). Jointly event extraction and visualization on Twitter via probabilistic modelling. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 7–12 August 2016, pp. 269–278.
- Zhou, Deyu, Xuan Zhang, Yin Zhou, Quan Zhao & Xin Geng (2016b). Emotion distribution learning from texts. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 638–647.
- Zhou, Dong, Séamus Lawless, Xuan Wu, Wenyu Zhao & Jianxun Liu (2016c). Enhanced personalized search using social data. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, Tex., 1–5 November 2016, pp. 700–710.
- Zhou, Lina, Judee K. Burgeon & Douglas P. Twitchell (2003). A longitudinal analysis of language behavior of deception in e-mail. In *Proceedings of the First NSF/NIJ Symposium on Intelligence and Security Informatics*, Tucson, Ariz., 2003, p. ??
- Zhou, Lina, Judee K. Burgeon, Douglas P. Twitchell, Tiantian Qin & Jr. Nunamaker, Jay F. (2004). A comparison of classification methods for predicting deception in computer-mediated communication. *Journal of Management Information Systems*, 20(4):139–165.
- Zhu, Tao, David Phipps, Adam Pridgen, Jedidiah R. Crandall & Dan S. Wallach (2013). The velocity of censorship: High-fidelity detection of microblog post deletions. In *Proceedings of the 22nd USENIX Security Symposium*, Washington, D.C., 14–16 August 2013.
- Zubiaga, Arkaitz, Elena Kochkina, Maria Liakata, Rob Procter & Michal Lukasik (2016). Stance classification in rumours as a sequential task exploiting the tree structure of social media conversations. In *Proceedings of the 26th International Conference on Computational Linguistics*, Osaka, Japan, 11–16 December 2016, pp. 2438–2448.
- Zukov Gregoric, Andrej, Zhiyuan Luo & Bartal Veyhe (2016). IBC-C: A dataset for armed conflict analysis. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Berlin, Germany, 7–12 August 2016, pp. 374–379.