

ACL 2019

**Social Media Mining for Health Applications (#SMM4H)
Workshop & Shared Task**

Proceedings of the Fourth Workshop

August 2, 2019
Florence, Italy

©2019 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-950737-46-8

Preface

Welcome to the 4th Social Media Mining for Health Applications Workshop and Shared Task - #SMM4H 2019.

The total number of users of social media continues to grow worldwide, resulting in the generation of vast amounts of data. Popular social networking sites such as Facebook, Twitter and Instagram dominate this sphere. According to estimates, 500 million tweets and 4.3 billion Facebook messages are posted every day ¹. The latest Pew Research Report ², nearly half of adults worldwide and two-thirds of all American adults (65%) use social networking. The report states that of the total users, 26% have discussed health information, and, of those, 30% changed behavior based on this information and 42% discussed current medical conditions. Advances in automated data processing, machine learning and NLP present the possibility of utilizing this massive data source for biomedical and public health applications, if researchers address the methodological challenges unique to this media.

In its fourth iteration, the #SMM4H workshop takes place in Florence, Italy, on August 2, 2019, and is co-located with the annual meeting of the Association of Computational Linguistic (ACL). Following on the success of our Workshops and accompanying shared tasks on the topic that were hosted at the Pacific Symposium in Biocomputing (PSB) in 2016, at the AMIA Annual Conference in 2017, and at the EMNLP conference in 2018, this workshop aims to provide a forum for the ACL community members to present and discuss NLP advances specific to social media use in the particularly challenging area of health applications, with a special focus given to automatic methods for the collection, extraction, representation, analysis, and validation of social media data for health informatics.

We received very high quality submissions for the workshop and selected only 8 articles for long presentations (Workshop Acceptance Rate: 54%) and 17 for short talks or posters presentations. As for the previous years, we ran in parallel to the workshop a shared task with a particular interest on social media mining for pharmacovigilance. For this fourth execution of the #SMM4H shared tasks, we challenged the community with two different problems involving annotated user posts from Twitter (tweets). The first problem focuses on performing pharmacovigilance from social media data during a series of three subtasks inviting the participants to extract and normalize tweets mentioning adverse effects of drugs. The second problem explores the generalizability of predictive models through a task of automatic classification of tweets with personal health experience mentions in multiple contexts. With a total of 34 teams registered and 19 teams having submitted a run, we confirm a growing interest of the community for health mining in social media data.

This year, we standardized the submission process during the shared task using the web platform Codalab³. We believed this helped improved the reproducibility of the experiments. Acting as a central hub, the data are easily distributed to the research community at large and, as the Codalab website remains active even after the competition, new teams can upload their submissions and be automatically evaluated to compare their results with the official results of the challenge.

Next year, we will go further to improve reproducibility by allowing participants to upload their code and models directly in Codalab for evaluation, a change which will guarantee a fair competition and the dissemination of the technical characteristics of the systems. Another important change this year that will impact future iterations of the workshop is an open call to the community for shared task proposals, ensuring that our workshop continues to address the main problems and challenges in this growing field

¹Team Gwava. "How Much Data is Created on the Internet each Day?" 2016, Available online at <https://www.gwava.com/blog/internet-data-created-daily>. [Accessed: 03-Jan-2017].

²Pew Research Center. "Social Media Fact Sheet". 2017. Available online at <http://www.pewinternet.org/fact-sheet/social-media/>. [Accessed: 03-Mar-2017].

³CodaLab is free and open-source, available at <https://competitions.codalab.org/>.

of health mining for social media.

The organizing committee would like to thank the program committee, consisting of 13 researchers, for their thoughtful input on the submissions, as well as the organizers of the ACL for their support and management. Finally, a huge thanks to all authors who submitted a paper to the workshop or participated in the shared tasks; this workshop would not have been possible without them and their hard work.

Graciela, Davy, Abeed, Arjun, Ashlynn, Michael

Organizing Committee:

Graciela Gonzalez-Hernandez, University of Pennsylvania
Davy Weissenbacher, University of Pennsylvania
Abeed Sarker, University of Pennsylvania
Michael Paul, University of Colorado-Boulder
Karen O'Connor, University of Pennsylvania
Ari Z. Klein, University of Pennsylvania
Arjun Magge, Arizona State University
Ashlynn R. Daughton, University of Colorado-Boulder

Program Committee:

Nigel Collier, University of Cambridge, UK
Larry Hunter, University of Colorado, USA
Hongfang Liu, Mayo Clinic Rochester, USA
Pierre Zweigenbaum, French National Center for Scientific Research, France
Cecile Paris, CSIRO, Australia
Kirk Roberts, University of Texas Houston, USA
Robert Leaman, US National Library of Medicine, USA
Azadeh Nikfarjam, Nuance Communications, USA
Ehsan Emazadeh, Google Inc., USA
Yutaka Sasaki, Toyota Technological Institute, Japan
Anne-Lyse Minard, Université d'Orléans, France
Thierry Poibeau, French National Center for Scientific Research, France
Kevin Cohen, University of Colorado, USA
Nicolas Turenne, French National Institute for Agricultural Research, France

Additional Reviewers:

Pramod Chandrashekar, Arizona State University, USA
Anahita Davoudi, University of Pennsylvania, USA

Table of Contents

<i>Extracting Kinship from Obituary to Enhance Electronic Health Records for Genetic Research</i> Kai He, Jialun Wu, Xiaoyong Ma, Chong Zhang, Ming Huang, Chen Li and Lixia Yao	1
<i>Lexical Normalization of User-Generated Medical Text</i> Anne Dirkson, Suzan Verberne and Wessel Kraaij	11
<i>Overview of the Fourth Social Media Mining for Health (SMM4H) Shared Tasks at ACL 2019</i> Davy Weissenbacher, Abeed Sarker, Arjun Magge, Ashlynn Daughton, Karen O’Connor, Michael J. Paul and Graciela Gonzalez-Hernandez	21
<i>MedNorm: A Corpus and Embeddings for Cross-terminology Medical Concept Normalisation</i> Maksim Belousov, William G. Dixon and Goran Nenadic	31
<i>Passive Diagnosis Incorporating the PHQ-4 for Depression and Anxiety</i> Fionn Delahunty, Robert Johansson and Mihael Arcan	40
<i>HITSZ-ICRC: A Report for SMM4H Shared Task 2019-Automatic Classification and Extraction of Adverse Effect Mentions in Tweets</i> Shuai Chen, Yuanhang Huang, Xiaowei Huang, Haoming Qin, Jun Yan and Buzhou Tang	47
<i>KFU NLP Team at SMM4H 2019 Tasks: Want to Extract Adverse Drugs Reactions from Tweets? BERT to The Rescue</i> Zulfat Miftahutdinov, Ilseyar Alimova and Elena Tutubalina	52
<i>Approaching SMM4H with Merged Models and Multi-task Learning</i> Tilia Ellendorff, Lenz Furrer, Nicola Colic, Noëmi Aepli and Fabio Rinaldi	58
<i>Identifying Adverse Drug Events Mentions in Tweets Using Attentive, Collocated, and Aggregated Medical Representation</i> Xinyan Zhao, Deahan Yu and V.G.Vinod Vydiswaran	62
<i>Correlating Twitter Language with Community-Level Health Outcomes</i> Arno Schneuwly, Ralf Grubenmann, Séverine Rion Logean, Mark Cieliebak and Martin Jaggi	71
<i>Affective Behaviour Analysis of On-line User Interactions: Are On-line Support Groups More Therapeutic than Twitter?</i> Giuliano Tortoreto, Evgeny Stepanov, Alessandra Cervone, Mateusz Dubiel and Giuseppe Riccardi	79
<i>Transfer Learning for Health-related Twitter Data</i> Anne Dirkson and Suzan Verberne	89
<i>NLP@UNED at SMM4H 2019: Neural Networks Applied to Automatic Classifications of Adverse Effects Mentions in Tweets</i> Javier Cortes-Tejada, Juan Martinez-Romo and Lourdes Araujo	93
<i>Detecting and Extracting of Adverse Drug Reaction Mentioning Tweets with Multi-Head Self Attention</i> Suyu Ge, Tao Qi, Chuhan Wu and Yongfeng Huang	96
<i>Deep Learning for Identification of Adverse Effect Mentions In Twitter Data</i> Paul Barry and Ozlem Uzuner	99

<i>Using Machine Learning and Deep Learning Methods to Find Mentions of Adverse Drug Reactions in Social Media</i>	
Pilar López Úbeda, Manuel Carlos Díaz Galiano, Maite Martin and L. Alfonso Urena Lopez . .	102
<i>Towards Text Processing Pipelines to Identify Adverse Drug Events-related Tweets: University of Michigan @ SMM4H 2019 Task 1</i>	
V.G.Vinod Vydiswaran, Grace Ganzel, Bryan Romas, Deahan Yu, Amy Austin, Neha Bhomia, Socheatha Chan, Stephanie Hall, Van Le, Aaron Miller, Olawunmi Oduyebo, Aulia Song, Radhika Sondhi, Danny Teng, Hao Tseng, Kim Vuong and Stephanie Zimmerman	107
<i>Neural Network to Identify Personal Health Experience Mention in Tweets Using BioBERT Embeddings</i>	
Shubham Gondane	110
<i>Give It a Shot: Few-shot Learning to Normalize ADR Mentions in Social Media Posts</i>	
Emmanouil Manousogiannis, Sepideh Mesbah, Alessandro Bozzon, Selene Baez and Robert Jan Sips	114
<i>BIGODM System in the Social Media Mining for Health Applications Shared Task 2019</i>	
Chen-Kai Wang, Hong-Jie Dai and Bo-Hung Wang	117
<i>Detection of Adverse Drug Reaction Mentions in Tweets Using ELMo</i>	
Sarah Sarabadani	120
<i>Adverse Drug Effect and Personalized Health Mentions, CLaC at SMM4H 2019, Tasks 1 and 4</i>	
Parsa Bagherzadeh, Nadia Sheikh and Sabine Bergler	123
<i>MIDAS@SMM4H-2019: Identifying Adverse Drug Reactions and Personal Health Experience Mentions from Twitter</i>	
Debanjan Mahata, Sarthak Anand, Haimin Zhang, Simra Shahid, Laiba Mehnaz, Yaman Kumar and Rajiv Ratn Shah	127
<i>Detection of Adverse Drug Reaction in Tweets Using a Combination of Heterogeneous Word Embeddings</i>	
Segun Taofeek Aroyehun and Alexander Gelbukh	133
<i>Identification of Adverse Drug Reaction Mentions in Tweets – SMM4H Shared Task 2019</i>	
Samarth Rawal, Siddharth Rawal, Saadat Anwar and Chitta Baral	136

Conference Program

Friday, August 2, 2019

- 9:00–9:10 *Introduction*
Graciela Gonzalez-Hernandez
- 9:10–9:30 *Extracting Kinship from Obituary to Enhance Electronic Health Records for Genetic Research*
Kai He, Jialun Wu, Xiaoyong Ma, Chong Zhang, Ming Huang, Chen Li and Lixia Yao
- 9:30–9:50 *Lexical Normalization of User-Generated Medical Text*
Anne Dirkson, Suzan Verberne and Wessel Kraaij
- 9:50–10:10 *Overview of the Fourth Social Media Mining for Health (SMM4H) Shared Tasks at ACL 2019*
Davy Weissenbacher, Abeed Sarker, Arjun Magge, Ashlynn Daughton, Karen O’Connor, Michael J. Paul and Graciela Gonzalez-Hernandez
- 10:30–11:30** *Coffee Break and Poster Session*
- 11:30–11:50 *MedNorm: A Corpus and Embeddings for Cross-terminology Medical Concept Normalisation*
Maksim Belousov, William G. Dixon and Goran Nenadic
- 11:50–12:00 *Passive Diagnosis Incorporating the PHQ-4 for Depression and Anxiety*
Fionn Delahunty, Robert Johansson and Mihael Arcan
- 12:00–14:00** *Lunch*
- 14:00–14:20 *HITSZ-ICRC: A Report for SMM4H Shared Task 2019-Automatic Classification and Extraction of Adverse Effect Mentions in Tweets*
Shuai Chen, Yuanhang Huang, Xiaowei Huang, Haoming Qin, Jun Yan and Buzhou Tang
- 14:20–14:40 *KFU NLP Team at SMM4H 2019 Tasks: Want to Extract Adverse Drugs Reactions from Tweets? BERT to The Rescue*
Zulfat Miftahutdinov, Ilseyar Alimova and Elena Tutubalina
- 14:40–15:00 *Approaching SMM4H with Merged Models and Multi-task Learning*
Tilia Ellendorff, Lenz Furrer, Nicola Colic, Noëmi Aepli and Fabio Rinaldi
- 15:00–16:00** *Coffee Break and Poster Session*

Friday, August 2, 2019 (continued)

- 16:00–16:20 *Identifying Adverse Drug Events Mentions in Tweets Using Attentive, Collocated, and Aggregated Medical Representation*
Xinyan Zhao, Deahan Yu and V.G.Vinod Vydiswaran
- 16:20–16:40 *Correlating Twitter Language with Community-Level Health Outcomes*
Arno Schneuwly, Ralf Grubenmann, Séverine Rion Logean, Mark Cieliebak and Martin Jaggi
- 16:40–16:50 *Affective Behaviour Analysis of On-line User Interactions: Are On-line Support Groups More Therapeutic than Twitter?*
Giuliano Tortoreto, Evgeny Stepanov, Alessandra Cervone, Mateusz Dubiel and Giuseppe Riccardi

17:00–17:30 *Discussion and Conclusion*

Poster Sessions

Transfer Learning for Health-related Twitter Data

Anne Dirkson and Suzan Verberne

NLP@UNED at SMM4H 2019: Neural Networks Applied to Automatic Classifications of Adverse Effects Mentions in Tweets

Javier Cortes-Tejada, Juan Martinez-Romo and Lourdes Araujo

Detecting and Extracting of Adverse Drug Reaction Mentioning Tweets with Multi-Head Self Attention

Suyu Ge, Tao Qi, Chuhan Wu and Yongfeng Huang

Deep Learning for Identification of Adverse Effect Mentions In Twitter Data

Paul Barry and Ozlem Uzuner

Using Machine Learning and Deep Learning Methods to Find Mentions of Adverse Drug Reactions in Social Media

Pilar López Úbeda, Manuel Carlos Díaz Galiano, Maite Martin and L. Alfonso Urena Lopez

Towards Text Processing Pipelines to Identify Adverse Drug Events-related Tweets: University of Michigan @ SMM4H 2019 Task 1

V.G.Vinod Vydiswaran, Grace Ganzel, Bryan Romas, Deahan Yu, Amy Austin, Neha Bhomia, Socheatha Chan, Stephanie Hall, Van Le, Aaron Miller, Olawunmi Oduyebo, Aulia Song, Radhika Sondhi, Danny Teng, Hao Tseng, Kim Vuong and Stephanie Zimmerman

Neural Network to Identify Personal Health Experience Mention in Tweets Using BioBERT Embeddings

Shubham Gondane

Friday, August 2, 2019 (continued)

Give It a Shot: Few-shot Learning to Normalize ADR Mentions in Social Media Posts

Emmanouil Manousogiannis, Sepideh Mesbah, Alessandro Bozzon, Selene Baez and Robert Jan Sips

BIGODM System in the Social Media Mining for Health Applications Shared Task 2019

Chen-Kai Wang, Hong-Jie Dai and Bo-Hung Wang

Detection of Adverse Drug Reaction Mentions in Tweets Using ELMo

Sarah Sarabadani

Adverse Drug Effect and Personalized Health Mentions, CLaC at SMM4H 2019, Tasks 1 and 4

Parsa Bagherzadeh, Nadia Sheikh and Sabine Bergler

MIDAS@SMM4H-2019: Identifying Adverse Drug Reactions and Personal Health Experience Mentions from Twitter

Debanjan Mahata, Sarthak Anand, Haimin Zhang, Simra Shahid, Laiba Mehnaz, Yaman Kumar and Rajiv Ratn Shah

Detection of Adverse Drug Reaction in Tweets Using a Combination of Heterogeneous Word Embeddings

Segun Taofeek Aroyehun and Alexander Gelbukh

Identification of Adverse Drug Reaction Mentions in Tweets – SMM4H Shared Task 2019

Samarth Rawal, Siddharth Rawal, Saadat Anwar and Chitta Baral

