BioNLP 2017

SIGBioMed Workshop on Biomedical Natural Language Processing

Proceedings of the 16th BioNLP Workshop

August 4, 2017
Vancouver, Canada

ISBN 978-1-945626-59-3

Biomedical natural language processing in 2017: The view from computational linguistics

Kevin Bretonnel Cohen, Dina Demner-Fushman, Sophia Ananiadou, and Jun-ichi Tsujii

According to the Association for Computational Linguistics guidelines on special interest groups (SIGs), *The function of a SIG is to encourage interest and activity in specific areas within the ACL's field*[1]. Is the SIGBioMed special interest group "within the ACL's field"? The titles of this year's papers suggest that it is, in that the current interest in deep learning in its many and varied manifestations is mirrored in those titles. Do those papers cover a specific area? They do, and in doing so, they demonstrate one of the great satisfactions of working in biomedical natural language processing.

One of the joys of involvement in the biomedical natural language processing community is seeing the development of research with clinical applications. As examples of such work being presented at BioNLP 2017, we would like to point out the two papers that discuss the application of natural language processing to the diagnosis of neurological disorders. Bhatia et al.[2] describe an approach to using speech processing in the assessment of patients with amyotrophic lateral sclerosis (also known as Lou Gehrig's disease), one of the more horrific motor neuron diseases. Good assessment of amyotrophic lateral sclerosis patients is important for a number of reasons, including the fact that accurate tracking of the inevitable deterioration that is a hallmark of this disease gives patients and their families the possibility of purposeful planning for the attendant disability and death. However, current methodologies for evaluating the status of amyotrophic lateral sclerosis patients necessarily involve expensive equipment and highly trained personnel; when further developed, this methodology could make such evaluation much more, and more frequently, available to ALS patients. The fact that the work reported here involves a speech modality is especially exciting, as speech-related indicators of future ALS can be present long before diagnosis. The paper uses measurements of phonological features of speech and their divergence from a baseline, and demonstrates correlation with physiological measures.

Adams et al.[3] describe work on detecting and categorizing word production errors associated with anomia, a particular kind of inability to find words. Screening for anomia is important because anomia is a symptom of stroke, but it is difficult and time-consuming to do, and therefore is not done as often as it should be. Automatic detection of anomia could be a nice enabler of improved care for stroke victims, but it is made difficult due to the subtlety of the phonological and semantic judgments that have to be made when assessing the phenomenon. The paper uses a combination of language modeling and phonologically-based edit distance calculation to approach the task, applying these techniques to data from the AphasiaBank collection of transcribed aphasic and healthy speech.

Although we have summarized only these two examples that address neurological disorders, there are several other papers on the use of natural language processing in clinical applications: patient-produced content in dementia [4], and health records ([5] on sepsis, [6] on e-cig use, [7] on pain and confusion); in the aggregate, these papers illustrate very nicely the potential for natural language processing to contribute to human well-being. Additionally, the current interest in the potential of natural language processing for social media is reflected in papers on studying medication intake via Twitter [8] and on monitoring dementia via blog posts [9]. Linguistics and language resources are represented in this year's papers, as well, including work on comparative structures [10] and a corpus construction effort [11].

The work in biomedical NLP was dominated by applications of deep learning to: punctuation restoration [12], text classification [13], relation extraction [14], [15], [16], information retrieval [17], and similarity judgments [18], among other exciting progress in biomedical language processing.

These are just a few examples of the high-quality research presented in BioNLP 2017.

In addition to the excellent submissions to the BioNLP workshop, this year features equally strong submissions to BioASQ challenge on large-scale biomedical semantic indexing and question answering, a shared task affiliated with BioNLP 2017. This year, the BioASQ challenge, which started in 2013, had three tasks:

- Large-Scale Online Biomedical Semantic Indexing
- Biomedical Semantic Question Answering
- Funding Information Extraction From Biomedical Literature

An overview of the tasks and the results of the challenge [19] are presented in an invited talk. The invited speaker, George Paliouras, is a senior researcher and head of the Intelligent Information Systems division of the Institute of Informatics and Telecommunications at NCSR "Demokritos", Greece. He holds a PhD in Machine Learning and has performed basic and applied research in Artificial Intelligence for the last 20 years. He is interested in the development of novel methods for addressing challenging big and small data analysis problems, such as learning complex models from structured relational data, learning from noisy and sparse data, learning from multiple heterogeneous data streams, and discovering patterns in hypergraphs. His research is motivated by the real-world problems. George has contributed to solving a variety of such problems, ranging from spam filtering and Web personalization to biomedical information retrieval. He has co-founded the spin-off company em i-sieve Technologies, which provides online reputation monitoring services.

Among various contributions to the research community, George Paliouras has served as board member in national and international scientific societies; he is serving on the editorial boards of international journals, and has chaired international conferences. He is involved in several research projects, in the role of scientific coordinator/principal investigator in some of them. In particular, he has coordinated and provided the infrastructure for the BioASQ project that was funded by the European Commission. He is currently coordinating iASiS, another project funded by the European Commission to develop big data integration and analysis methods that will provide insight to public health policy-making for personalized medicine.

Acknowledging the community

As always, the organizers thank the authors who submitted their work to BioNLP 2017 —without them, there would be no meeting, no opportunity to share the progress and the pain of the past year with the community. We have listed above only a few of the exceptional submissions that were accepted for oral (20) and poster (28) presentations.

The distribution of scores this year suggests that a large amount of excellent work was submitted for review and resulted in 77% acceptance ratio. At the same time, the distribution suggests that the reviewers were careful and thorough, and the organizers thank them for that, and for thoroughly reviewing up to five papers on a very tight schedule.

We greatly appreciate the BioNLP core authors and program committee members who have been building up the community and the workshop for the past sixteen years. We are also happy to see the excellent new submissions and the new reviewers, and hope they will continue contributing to BioNLP.

References

[1] SIG Creation Guidelines *https://goo.gl/yIQCHo* 2017.

[2] Archna Bhatia, Bonnie Dorr, Kristy Hollingshead, Samuel L. Phillips and Barbara McKenzie *Characterization of Divergence in Impaired Speech of ALS Patients* 2017.

[3] Joel Adams, Steven Bedrick, Gerasimos Fergadiotis, Kyle Gorman and Jan van Santen *Target word prediction and paraphasia classification in spoken discourse* 2017.

[4] Vaden Masrani, Gabriel Murray, Thalia Field and Giuseppe Carenini *Detecting Dementia through Retrospective Analysis of Routine Blog Posts by Bloggers with Dementia* 2017.

[5] Emilia Apostolova and Tom Velez *Toward Automated Early Sepsis Alerting: Identifying Infection Patients from Nursing Notes* 2017.

[6] Danielle Mowery, Brett South, Olga Patterson, Shu-Hong Zhu and Mike Conway *Investigating the Documentation of Electronic Cigarette Use in the Veteran Affairs Electronic Health Record: A Pilot Study* 2017.

[7] Hans Moen, Kai Hakala, Farrokh Mehryary, Laura-Maria Peltonen, Tapio Salakoski, Filip Ginter and Sanna Salanterä *Detecting mentions of pain and acute confusion in Finnish clinical text* 2017.

[8] Ari Klein, Abeed Sarker, Masoud Rouhizadeh, Karen O'Connor and Graciela Gonzalez *Detecting Personal Medication Intake in Twitter: An Annotated Corpus and Baseline Classification System* 2017.

[9] Vaden Masrani, Gabriel Murray, Thalia Field and Giuseppe Carenini *Detecting Dementia through Retrospective Analysis of Routine Blog Posts by Bloggers with Dementia* 2017.

[10] Samir Gupta, A.S.M. Ashique Mahmood, Karen Ross, Cathy Wu and K. Vijay-Shanker *Identifying Comparative Structures in Biomedical Text* 2017.

[11] Rezarta Islamaj Dogan, Andrew Chatr-aryamontri, Sun Kim, Chih-Hsuan Wei, Yifan Peng, Donald Comeau and Zhiyong Lu *BioCreative VI Precision Medicine Track: creating a training corpus for mining protein-protein interactions affected by mutations* 2017.

[12] Wael Salloum, Greg Finley, Erik Edwards, Mark Miller and David Suendermann-Oeft *Deep Learning for Punctuation Restoration in Medical Reports* 2017.

[13] Simon Baker and Anna Korhonen *Initializing neural networks for hierarchical multi-label text classification* 2017.

[14] Chen Lin, Timothy Miller, Dmitriy Dligach, Steven Bethard and Guergana Savova *Representations of Time Expressions for Temporal Relation Extraction with Convolutional Neural Networks* 2017.

[15] Masaki Asada, Makoto Miwa and Yutaka Sasaki *Extracting Drug-Drug Interactions with Attention CNNs* 2017.

[16] Yifan Peng and Zhiyong Lu *Deep learning for extracting protein-protein interactions from biomedical literature* 2017.

[17] Sunil Mohan, Nicolas Fiorini, Sun Kim and Zhiyong Lu *Deep Learning for Biomedical Information Retrieval: Learning Textual Relevance from Click Logs* 2017.

[18] Bridget McInnes and Ted Pedersen *Improving Correlation with Human Judgments by Integrating Semantic Similarity with Second–Order Vectors* 2017.

[19] Anastasios Nentidis, Konstantinos Bougiatiotis, Anastasia Krithara, Georgios Paliouras and Ioannis Kakadiaris *Results of the fifth edition of the BioASQ Challenge* 2017.

Organizers:

Kevin Bretonnel Cohen, University of Colorado School of Medicine, USA
Dina Demner-Fushman, US National Library of Medicine
Sophia Ananiadou, National Centre for Text Mining and University of Manchester, UK
Jun-ichi Tsujii, National Institute of Advanced Industrial Science and Technology, Japan

Program Committee:

Sophia Ananiadou, National Centre for Text Mining and University of Manchester, UK
Ion Androutsopoulos, Athens University of Economics and Business, Greece
Emilia Apostolova, Language.ai, USA
Eiji Aramaki, University of Tokyo, Japan
Alan Aronson, US National Library of Medicine
Asma Ben Abacha, US National Library of Medicine
Olivier Bodenreider, US National Library of Medicine
Kevin Bretonnel Cohen, University of Colorado School of Medicine, USA
Leonardo Campillos Llanos, LIMSI - CNRS, France
Kevin Bretonnel Cohen, University of Colorado School of Medicine, USA
Nigel Collier, University of Cambridge, UK
Dina Demner-Fushman, US National Library of Medicine
Filip Ginter, University of Turku, Finland
Graciela Gonzalez, University of Pennsylvania, USA
Cyril Grouin, LIMSI - CNRS, France
Antonio Jimeno Yepes, IBM, Melbourne Area, Australia
Halil Kilicoglu, US National Library of Medicine
Aris Kosmopoulos, NCSR Demokritos, Greece
Robert Leaman, US National Library of Medicine
Chris Lu, US National Library of Medicine
Zhiyong Lu, US National Library of Medicine
Juan Miguel Cejuela, Technische Universität München, Germany
Timothy Miller, Children's Hospital Boston, USA
Makoto Miwa, Toyota Technological Institute, Japan
Danielle L Mowery, VA Salt Lake City Health Care System, USA
Diego Molla, Macquarie University, Australia
Jim Mork, National Library of Medicine, USA
Yassine Mrabet, US National Library of Medicine
Henning Müller, University of Applied Sciences, Switzerland
Claire Nédellec, INRA, France
Anastasios Nentidis, NCSR Demokritos, Athens, Greece
Aurélie Névéol, LIMSI - CNRS, France
Mariana Neves, Hasso Plattner Institute and University of Potsdam, Germany
Nhung Nguyen, The University of Manchester, Manchester
Naoaki Okazaki, Tohoku University, Japan
Georgios Paliouras, NCSR Demokritos, Athens, Greece
Ioannis Partalas, Viseo group, France
John Prager, Thomas J. Watson Research Center, IBM, USA
Sampo Pyysalo, University of Cambridge, UK

Francisco J. Ribadas-Pena, University of Vigo, Spain
Fabio Rinaldi, University of Zurich, Switzerland
Angus Roberts, The University of Sheffield, UK
Kirk Roberts, The University of Texas Health Science Center at Houston, USA
Hagit Shatkay, University of Delaware, USA
Pontus Stenetorp, University College London, UK
Karin Verspoor, The University of Melbourne, Australia
Ellen Voorhees, National Institute of Standards and Technology, USA
Byron C. Wallace, University of Texas at Austin, USA
W John Wilbur, US National Library of Medicine
Hai Zhao, Shanghai Jiao Tong University, Shanghai
Pierre Zweigenbaum, LIMSI - CNRS, France

Additional Reviewers:

Moumita Bhattacharya, University of Delaware, USA
Louise Deléger, INRA - MaIAGE, France
Lenz Furrer, Institute of Computational Linguistics, UZH, Zurich, Switzerland
Genevieve Gorrell, Sheffield University, UK
Ari Klein, University of Pennsylvania School of Medicine
Yifan Peng, US National Library of Medicine
Vassiliki Rentoumi, National Centre for Scientific Research Demokritos, Athens, Greece
Masoud Rouhizadeh, University of Pennsylvania, USA
Abeed Sarker, University of Pennsylvania, USA
Xingyi Song, University of Sheffield, UK
Tasnia Tahsin, Arizona State University, USA
Hegler Tissot, Federal University of Parana, Brazil
Ken Yano, Nara Institute of Science and Technology, Japan

Invited Speaker:

George Paliouras, National Centre for Scientific Research Demokritos, Athens, Greece

Table of Contents

Conference Program

Friday August 4, 2017

8:30–8:45 Opening remarks

8:45–10:30 Session 1: Prediction and relation extraction

8:45–9:00 *Target word prediction and paraphasia classification in spoken discourse*
Joel Adams, Steven Bedrick, Gerasimos Fergadiotis, Kyle Gorman and Jan van Santen

9:00–9:15 *Extracting Drug-Drug Interactions with Attention CNNs*
Masaki Asada, Makoto Miwa and Yutaka Sasaki

9:15–9:30 *Insights into Analogy Completion from the Biomedical Domain*
Denis Newman-Griffis, Albert Lai and Eric Fosler-Lussier

9:30–9:45 *Deep learning for extracting protein-protein interactions from biomedical literature*
Yifan Peng and Zhiyong Lu

9:45–10:00 *Stacking With Auxiliary Features for Entity Linking in the Medical Domain*
Nazneen Fatema Rajani, Mihaela Bornea and Ken Barker

10:00–10:30 Invited Talk: "Results of the 5th edition of BioASQ Challenge" – Georgios Paliouras

Results of the fifth edition of the BioASQ Challenge
Anastasios Nentidis, Konstantinos Bougiatiotis, Anastasia Krithara, Georgios Paliouras and Ioannis Kakadiaris

10:30–11:00 ***Coffee Break***

Friday August 4, 2017 (continued)

11:00–12:30 Session 2: BioASQ 2017 and more

11:00–11:15 *Tackling Biomedical Text Summarization: OAQA at BioASQ 5B*
Khyathi Chandu, Aakanksha Naik, Aditya Chandrasekar, Zi Yang, Niloy Gupta and Eric Nyberg

11:15–11:30 *Macquarie University at BioASQ 5b – Query-based Summarisation Techniques for Selecting the Ideal Answers*
Diego Molla

11:30–11:45 *Neural Question Answering at BioASQ 5B*
Georg Wiese, Dirk Weissenborn and Mariana Neves

11:45–12:00 *End-to-End System for Bacteria Habitat Extraction*
Farrokh Mehryary, Kai Hakala, Suwisa Kaewphan, Jari Björne, Tapio Salakoski and Filip Ginter

12:00–12:15 *Creation and evaluation of a dictionary-based tagger for virus species and proteins*
Helen Cook, Rudolfs Berzins, Cristina Leal Rodrıguez, Juan Miguel Cejuela and Lars Juhl Jensen

12:15–12:30 *Representation of complex terms in a vector space structured by an ontology for a normalization task*
Arnaud Ferré, Pierre Zweigenbaum and Claire Nédellec

12:30–14:00 ***Lunch break***

Friday August 4, 2017 (continued)

14:00–15:30 Session 3: From bio to clinical NLP

14:00–14:15 *Improving Correlation with Human Judgments by Integrating Semantic Similarity with Second–Order Vectors*
Bridget McInnes and Ted Pedersen

14:15–14:30 *Proactive Learning for Named Entity Recognition*
Maolin Li, Nhung Nguyen and Sophia Ananiadou

14:30–14:45 *Biomedical Event Extraction using Abstract Meaning Representation*
Sudha Rao, Daniel Marcu, Kevin Knight and Hal Daumé III

14:45–15:00 *Detecting Personal Medication Intake in Twitter: An Annotated Corpus and Baseline Classification System*
Ari Klein, Abeed Sarker, Masoud Rouhizadeh, Karen O'Connor and Graciela Gonzalez

15:00–15:15 *Unsupervised Context-Sensitive Spelling Correction of Clinical Free-Text with Word and Character N-Gram Embeddings*
Pieter Fivez, Simon Suster and Walter Daelemans

15:15–15:30 *Characterization of Divergence in Impaired Speech of ALS Patients*
Archna Bhatia, Bonnie Dorr, Kristy Hollingshead, Samuel L. Phillips and Barbara McKenzie

15:30–16:00 ***Coffee Break***

Friday August 4, 2017 (continued)

16:00–16:30 Session 4 More clinical NLP

16:00–16:15 *Deep Learning for Punctuation Restoration in Medical Reports*
Wael Salloum, Greg Finley, Erik Edwards, Mark Miller and David Suendermann-Oeft

16:15–16:30 *Unsupervised Domain Adaptation for Clinical Negation Detection*
Timothy Miller, Steven Bethard, Hadi Amiri and Guergana Savova

16:30–18:00 Poster Session

BioCreative VI Precision Medicine Track: creating a training corpus for mining protein-protein interactions affected by mutations
Rezarta Islamaj Dogan, Andrew Chatr-aryamontri, Sun Kim, Chih-Hsuan Wei, Yifan Peng, Donald Comeau and Zhiyong Lu

Painless Relation Extraction with Kindred
Jake Lever and Steven Jones

Noise Reduction Methods for Distantly Supervised Biomedical Relation Extraction
Gang Li, Cathy Wu and K. Vijay-Shanker

Role-Preserving Redaction of Medical Records to Enable Ontology-Driven Processing
Seth Polsley, Atif Tahir, Muppala Raju, Akintayo Akinleye and Duane Steward

Annotation of pain and anesthesia events for surgery-related processes and outcomes extraction
Wen-wai Yim, Dario Tedesco, Catherine Curtin and Tina Hernandez-Boussard

Identifying Comparative Structures in Biomedical Text
Samir Gupta, A.S.M. Ashique Mahmood, Karen Ross, Cathy Wu and K. Vijay-Shanker

Tagging Funding Agencies and Grants in Scientific Articles using Sequential Learning Models
Subhradeep Kayal, Zubair Afzal, George Tsatsaronis, Sophia Katrenko, Pascal Coupet, Marius Doornenbal and Michelle Gregory

Deep Learning for Biomedical Information Retrieval: Learning Textual Relevance from Click Logs
Sunil Mohan, Nicolas Fiorini, Sun Kim and Zhiyong Lu

Friday August 4, 2017 (continued)

Detecting Dementia through Retrospective Analysis of Routine Blog Posts by Bloggers with Dementia
Vaden Masrani, Gabriel Murray, Thalia Field and Giuseppe Carenini

Protein Word Detection using Text Segmentation Techniques
Devi Ganesan, Ashish V. Tendulkar and Sutanu Chakraborti

External Evaluation of Event Extraction Classifiers for Automatic Pathway Curation: An extended study of the mTOR pathway
Wojciech Kusa and Michael Spranger

Toward Automated Early Sepsis Alerting: Identifying Infection Patients from Nursing Notes
Emilia Apostolova and Tom Velez

Enhancing Automatic ICD-9-CM Code Assignment for Medical Texts with PubMed
Danchen Zhang, Daqing He, Sanqiang Zhao and Lei Li

Evaluating Feature Extraction Methods for Knowledge-based Biomedical Word Sense Disambiguation
Sam Henry, Clint Cuffy and Bridget McInnes

Investigating the Documentation of Electronic Cigarette Use in the Veteran Affairs Electronic Health Record: A Pilot Study
Danielle Mowery, Brett South, Olga Patterson, Shu-Hong Zhu and Mike Conway

Automated Preamble Detection in Dictated Medical Reports
Wael Salloum, Greg Finley, Erik Edwards, Mark Miller and David Suendermann-Oeft

A Biomedical Question Answering System in BioASQ 2017
Mourad Sarrouti and Said Ouatik El Alaoui

Adapting Pre-trained Word Embeddings For Use In Medical Coding
Kevin Patel, Divya Patel, Mansi Golakiya, Pushpak Bhattacharyya and Nilesh Birari

Initializing neural networks for hierarchical multi-label text classification
Simon Baker and Anna Korhonen

Biomedical Event Trigger Identification Using Bidirectional Recurrent Neural Network Based Models
Rahul V S S Patchigolla, Sunil Sahu and Ashish Anand

Representations of Time Expressions for Temporal Relation Extraction with Convolutional Neural Networks
Chen Lin, Timothy Miller, Dmitriy Dligach, Steven Bethard and Guergana Savova

Automatic Diagnosis Coding of Radiology Reports: A Comparison of Deep Learning and Conventional Classification Methods
Sarvnaz Karimi, Xiang Dai, Hamedh Hassanzadeh and Anthony Nguyen

Automatic classification of doctor-patient questions for a virtual patient record query task
Leonardo Campillos Llanos, Sophie Rosset and Pierre Zweigenbaum

Assessing the performance of Olelo, a real-time biomedical question answering application
Mariana Neves, Fabian Eckert, Hendrik Folkerts and Matthias Uflacker

Clinical Event Detection with Hybrid Neural Architecture
Adyasha Maharana and Meliha Yetisgen

Extracting Personal Medical Events for User Timeline Construction using Minimal Supervision
Aakanksha Naik, Chris Bogart and Carolyn Rose

Detecting mentions of pain and acute confusion in Finnish clinical text
Hans Moen, Kai Hakala, Farrokh Mehryary, Laura-Maria Peltonen, Tapio Salakoski, Filip Ginter and Sanna Salanterä

A Multi-strategy Query Processing Approach for Biomedical Question Answering: USTB_PRIR at BioASQ 2017 Task 5B
Zan-Xia Jin, Bo-Wen Zhang, Fan Fang, Le-Le Zhang and Xu-Cheng Yin