

EMNLP 2018



Fact Extraction and VERification

Proceedings of the First Workshop

November 1, 2018
Brussels, Belgium

We thank our sponsor Amazon Research in Cambridge for their generous support.



©2018 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 978-1-948087-73-5

Introduction

With billions of individual pages on the web providing information on almost every conceivable topic, we should have the ability to collect facts that answer almost every conceivable question. However, only a small fraction of this information is contained in structured sources (Wikidata, Freebase, etc.) – we are therefore limited by our ability to transform free-form text to structured knowledge. There is, however, another problem that has become the focus of a lot of recent research and media coverage: false information coming from unreliable sources.

In an effort to jointly address both problems, herein we proposed this workshop to promote research in joint Fact Extraction and VERification (FEVER). We aim for FEVER to be a long-term venue for work in verifiable knowledge extraction.

To stimulate progress in this direction, we also hosted the FEVER Challenge, an information verification shared task on a purposely-constructed dataset. We received entries from 23 competing teams, 19 of which scored higher than the previously published baseline. We invited descriptions of the participating systems and we received 15 system descriptions, all of which are included in these proceedings. We offered the top 4 systems oral presentations.

For the main workshop, we received 23 submissions, out of which we accepted 14 (3 oral presentations and 11 posters).

Organizers:

James Thorne (University of Sheffield)
Andreas Vlachos (University of Sheffield)
Oana Cocarascu (Imperial College London)
Christos Christodoulopoulos (Amazon)
Arpit Mittal (Amazon)

Program Committee:

Nikolaos Aletras (University of Sheffield), Fernando Alva-Manchego (University of Sheffield), Isabelle Augenstein (University of Copenhagen), Esma Balkir (University of Edinburgh), Daniele Bonadiman (University of Trento), Matko Bošnjak (University College London), Kris Cao (University of Cambridge), Tuhin Chakrabarty (Columbia University), Weiwei Cheng (Amazon), Bich-Ngoc Do (Heidelberg University), Micha Elsner (The Ohio State University), Diego Esteves (Universität Bonn), Frédéric Godin (ELIS - IDLab, Ghent University), Ivan Habernal (UKP Lab, Technische Universität Darmstadt), Andreas Hanselowski (UKP lab, Technische Universität Darmstadt), Christopher Hidey (Columbia University), Julia Hockenmaier (University of Illinois Urbana-Champaign), Alexandre Klementiev (Amazon Development Center Germany), Jan Kowollik (University of Duisburg-Essen), Anjishnu Kumar (Amazon), Nayeon Lee (Hong Kong University of Science and Technology), Pranava Swaroop Madhyastha (University of Sheffield), Christopher Malon (NEC Laboratories America), Marie-Catherine de Marneffe (The Ohio State University), Stephen Mayhew (University of Pennsylvania), Marie-Francine Moens (KU Leuven), Jason Naradowsky (University College London), Yixin Nie (UNC), Farhad Nooralahzadeh (University of Oslo), Wolfgang Otto (GESIS – Leibniz-Institute for the Social Sciences in Cologne), Ankur Padia (University of Maryland, Baltimore County), Mithun Paul (University Of Arizona), Tamara Polajnar (University of Cambridge), Hoifung Poon (Microsoft Research), Preethi Raghavan (IBM Research TJ Watson), Marek Rei (University of Cambridge), Laura Rimell (DeepMind), Tim Rocktäschel (University College London and Facebook AI Research), Jodi Schneider (UIUC), Claudia Schulz (UKP Lab, Technische Universität Darmstadt), Diarmuid Ó Séaghdha (Apple), Sameer Singh (University of California, Irvine), Kevin Small (Amazon), Christian Stab (UKP Lab, Technische Universität Darmstadt), Motoki Taniguchi (Fuji Xerox), Paolo Torrioni (Alma Mater - Università di Bologna), Serena Villata (Université Côte d’Azur, CNRS, Inria, I3S), Zeerak Waseem (University of Sheffield), Noah Weber (Stony Brook University), Johannes Welbl (University College London), Menglin Xia (University of Cambridge), Takuma Yoneda (Toyota Technological Institute)

Invited Speakers:

Luna Dong (Amazon)
Marie-Francine Moens (KU Leuven)
Delip Rao (Joostware AI Research, Johns Hopkins University)
Tim Rocktäschel (Facebook AI Research, University College London)

Table of Contents

<i>The Fact Extraction and VERification (FEVER) Shared Task</i> James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos and Arpit Mittal	1
<i>The Data Challenge in Misinformation Detection: Source Reputation vs. Content Veracity</i> Fatemeh Torabi Asr and Maite Taboada	10
<i>Crowdsourcing Semantic Label Propagation in Relation Classification</i> Anca Dumitrache, Lora Aroyo and Chris Welty	16
<i>Retrieve and Re-rank: A Simple and Effective IR Approach to Simple Question Answering over Knowledge Graphs</i> Vishal Gupta, Manoj Chinnakotla and Manish Shrivastava	22
<i>Information Nutrition Labels: A Plugin for Online News Evaluation</i> Vincentius Kevin, Birte Högden, Claudia Schwenger, Ali Sahan, Neelu Madan, Piush Aggarwal, Anusha Bangaru, Farid Muradov and Ahmet Aker	28
<i>Joint Modeling for Query Expansion and Information Extraction with Reinforcement Learning</i> Motoki Taniguchi, Yasuhide Miura and Tomoko Ohkuma	34
<i>Towards Automatic Fake News Detection: Cross-Level Stance Detection in News Articles</i> Costanza Conforti, Mohammad Taher Pilehvar and Nigel Collier	40
<i>Belittling the Source: Trustworthiness Indicators to Obfuscate Fake News on the Web</i> Diego Esteves, Aniketh Janardhan Reddy, Piyush Chawla and Jens Lehmann	50
<i>Automated Fact-Checking of Claims in Argumentative Parliamentary Debates</i> Nona Naderi and Graeme Hirst	60
<i>Stance Detection in Fake News A Combined Feature Representation</i> Bilal Ghanem, Paolo Rosso and Francisco Rangel	66
<i>Zero-shot Relation Classification as Textual Entailment</i> Abiola Obamuyide and Andreas Vlachos	72
<i>Teaching Syntax by Adversarial Distraction</i> Juho Kim, Christopher Malon and Asim Kadav	79
<i>Where is Your Evidence: Improving Fact-checking by Justification Modeling</i> Tariq Alhindi, Savvas Petridis and Smaranda Muresan	85
<i>Affordance Extraction and Inference based on Semantic Role Labeling</i> Daniel Loureiro and Alípio Jorge	91
<i>UCL Machine Reading Group: Four Factor Framework For Fact Finding (HexaF)</i> Takuma Yoneda, Jeff Mitchell, Johannes Welbl, Pontus Stenetorp and Sebastian Riedel	97
<i>Multi-Sentence Textual Entailment for Claim Verification</i> Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz and Iryna Gurevych	103
<i>Team Papelo: Transformer Networks at FEVER</i> Christopher Malon	109

<i>Uni-DUE Student Team: Tackling fact checking through decomposable attention neural network</i> Jan Kowollik and Ahmet Aker	114
<i>SIRIUS-LTG: An Entity Linking Approach to Fact Extraction and Verification</i> Farhad Nooralahzadeh and Lilja Øvrelid	119
<i>Integrating Entity Linking and Evidence Ranking for Fact Extraction and Verification</i> Motoki Taniguchi, Tomoki Taniguchi, Takumi Takahashi, Yasuhide Miura and Tomoko Ohkuma 124	
<i>Robust Document Retrieval and Individual Evidence Modeling for Fact Extraction and Verification.</i> Tuhin Chakrabarty, Tariq Alhindi and Smaranda Muresan	127
<i>DeFactoNLP: Fact Verification using Entity Recognition, TFIDF Vector Comparison and Decomposable Attention</i> Aniketh Janardhan Reddy, Gil Rocha and Diego Esteves	132
<i>An End-to-End Multi-task Learning Model for Fact Checking</i> sizhen li, Shuai Zhao, Bo Cheng and Hao Yang	138
<i>Team GESIS Cologne: An all in all sentence-based approach for FEVER</i> Wolfgang Otto	145
<i>Team SWEEPer: Joint Sentence Extraction and Fact Checking with Pointer Networks</i> Christopher Hidey and Mona Diab	150
<i>QED: A fact verification system for the FEVER shared task</i> Jackson Luken, Nanjiang Jiang and Marie-Catherine de Marneffe	156
<i>Team UMBC-FEVER : Claim verification using Semantic Lexical Resources</i> Ankur Padia, Francis Ferraro and Tim Finin	161
<i>A mostly unlexicalized model for recognizing textual entailment</i> Mithun Paul, Rebecca Sharp and Mihai Surdeanu	166

Conference Program

Thursday, November 1, 2018

09:00–09:15 *Welcome Talk*
Organizers

09:15–10:00 *Invited Talk: Learning with Explanations*
Tim Rocktäschel

10:00–10:30 Research Talks 1

10:00–10:15 *The Data Challenge in Misinformation Detection: Source Reputation vs. Content Veracity*
Fatemeh Torabi Asr and Maite Taboada

10:15–10:30 *Towards Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection*
Lev Konstantinovskiy, Oliver Price, Mevan Babakar and Arkaitz Zubiaga

10:30–11:30 Research Posters

Crowdsourcing Semantic Label Propagation in Relation Classification
Anca Dumitrache, Lora Aroyo and Chris Welty

Retrieve and Re-rank: A Simple and Effective IR Approach to Simple Question Answering over Knowledge Graphs
Vishal Gupta, Manoj Chinnakotla and Manish Shrivastava

Information Nutrition Labels: A Plugin for Online News Evaluation
Vincentius Kevin, Birte Högden, Claudia Schwenger, Ali Sahan, Neelu Madan, Piush Aggarwal, Anusha Bangaru, Farid Muradov and Ahmet Aker

Joint Modeling for Query Expansion and Information Extraction with Reinforcement Learning
Motoki Taniguchi, Yasuhide Miura and Tomoko Ohkuma

Towards Automatic Fake News Detection: Cross-Level Stance Detection in News Articles
Costanza Conforti, Mohammad Taher Pilehvar and Nigel Collier

Belittling the Source: Trustworthiness Indicators to Obfuscate Fake News on the Web
Diego Esteves, Aniketh Janardhan Reddy, Piyush Chawla and Jens Lehmann

Thursday, November 1, 2018 (continued)

Automated Fact-Checking of Claims in Argumentative Parliamentary Debates

Nona Naderi and Graeme Hirst

Stance Detection in Fake News A Combined Feature Representation

Bilal Ghanem, Paolo Rosso and Francisco Rangel

Zero-shot Relation Classification as Textual Entailment

Abiola Obamuyide and Andreas Vlachos

Teaching Syntax by Adversarial Distraction

Juho Kim, Christopher Malon and Asim Kadav

Where is Your Evidence: Improving Fact-checking by Justification Modeling

Tariq Alhindi, Savvas Petridis and Smaranda Muresan

11:30–12:15 *Invited Talk: Argumentation Mining and Generation Supporting the Verification of Content*

Marie-Francine Moens

12:15–12:30 Research Talks 2

12:15–12:30 *Affordance Extraction and Inference based on Semantic Role Labeling*

Daniel Loureiro and Alípio Jorge

14:00–14:45 *Invited Talk: Building a broad knowledge graph for products*

Luna Dong

Thursday, November 1, 2018 (continued)

14:45–15:30 Shared Task Flash Talks

14:45–14:50 *The Fact Extraction and VERification (FEVER) Shared Task*

Organizers

14:50–15:00 *Combining Fact Extraction and Claim Verification in an NLI Model*

Yixin Nie, Haonan Chen and Mohit Bansal

15:00–15:10 *UCL Machine Reading Group: Four Factor Framework For Fact Finding (HexaF)*

Takuma Yoneda, Jeff Mitchell, Johannes Welbl, Pontus Stenetorp and Sebastian Riedel

15:10–15:20 *Multi-Sentence Textual Entailment for Claim Verification*

Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz and Iryna Gurevych

15:20–15:30 *Team Papelo: Transformer Networks at FEVER*

Christopher Malon

15:30–16:30 Shared Task Posters

Uni-DUE Student Team: Tackling fact checking through decomposable attention neural network

Jan Kowollik and Ahmet Aker

SIRIUS-LTG: An Entity Linking Approach to Fact Extraction and Verification

Farhad Nooralahzadeh and Lilja Øvrelid

Integrating Entity Linking and Evidence Ranking for Fact Extraction and Verification

Motoki Taniguchi, Tomoki Taniguchi, Takumi Takahashi, Yasuhide Miura and Tomoko Ohkuma

Robust Document Retrieval and Individual Evidence Modeling for Fact Extraction and Verification.

Tuhin Chakrabarty, Tariq Alhindi and Smaranda Muresan

DeFactoNLP: Fact Verification using Entity Recognition, TFIDF Vector Comparison and Decomposable Attention

Aniketh Janardhan Reddy, Gil Rocha and Diego Esteves

Thursday, November 1, 2018 (continued)

An End-to-End Multi-task Learning Model for Fact Checking

sizhen li, Shuai Zhao, Bo Cheng and Hao Yang

Team GESIS Cologne: An all in all sentence-based approach for FEVER

Wolfgang Otto

Team SWEEPer: Joint Sentence Extraction and Fact Checking with Pointer Networks

Christopher Hidey and Mona Diab

QED: A fact verification system for the FEVER shared task

Jackson Luken, Nanjiang Jiang and Marie-Catherine de Marneffe

Team UMBC-FEVER : Claim verification using Semantic Lexical Resources

Ankur Padia, Francis Ferraro and Tim Finin

A mostly unlexicalized model for recognizing textual entailment

Mithun Paul, Rebecca Sharp and Mihai Surdeanu

16:30–17:15 *Invited Talk: Call for Help: Putting Computation in Computational Fact Checking*

Delip Rao

17:15–17:30 *Prizes + Closing Remarks*

Organizers